An Interop Publication

ConneXions®

The Interoperability Report

July 1991 | *Special Issue: The Changing Face of the Internet* | Volume 5, No. 7

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.

ISSN 0894–5926

From the Editor

The Internet has grown up. It is no longer just a place where academics and researchers meet to exchange ideas. The Internet is everywhere, from small building LANs to regional, national and international networks. It is growing fast and employing new technology every day. The success of the TCP/IP protocol suite is clearly evident, just look at the ads and articles in any computer trade publication. This means that Internet technology is being used by far more than the estimated 3 million users of the connected Internet.

In this special issue, we take a step away from our technical perspective and look at some Internet developments. Our focus will be *applications:* the Internet is now being used for far more than the original "Big Three" applications (e-mail, file transfer and remote login).

We begin with an article by John Quarterman entitled "Networks: From Technology to Community." John describes how faster networks lead to new services, then new uses, and finally communities. The article is reprinted from the May 1991 issue of *Matrix News.*

This is followed by a report from Europe where the networking situation is "chaotic" to put it mildly. Bernhard Stockman describes efforts underway to coordinate the various networking projects with the aim of improving services nationally and internationally.

Library online catalog systems have become available over the Internet. Nearly 200 libraries are currently accessible, and this number is growing at the rate of more than one new system per week. Billy Baron gives an overview of current and future use of the Internet library computer systems and the information available on them.

The Internet is also becoming "more commercial" in the sense that various service providers now offer Internet connectivity. In March, three such providers announced an agreement to link their networks via a *Commercial Internet Exchange* (CIX). Daniel Dern reports on this development in an article on page 20.

As the nature of the Internet changes, the issue of "who pays" and "how do we charge" arises. This topic, generally referred to as *Internet Accounting,* is discussed by Don Hirsh starting on page 24.

Certain processing intensive applications (such as medical imaging) are creating a demand for very high speed networks. New technologies such as SONET, ATM, and HIPPI are making such networks a reality. We will discuss these technologies in future issues of *ConneXions.* This month, Carl Malamud describes a couple of high-speed networking projects from the perspective of "why do we need gigabit networks?" The article is adapted from his soon-to-be-published book, *STACKS—Interoperability in Today's Computer Networks.*

Networks: From Technology to Community

by John S. Quarterman

Networks are not just technology. Faster networks lead to new services, then new uses, then communities. This article discusses aspects of the history and near future of a few networks (the ARPANET, the Internet, UUCP, USENET, and BITNET) to illustrate some patterns that have occurred on many networks. The network services mentioned here are intended to be typical. If your favorite service is omitted, I'd be interested in hearing how you think it fits (or doesn't fit) in this scheme.

The technological information in the first three quarters of the article is actually background to the sketch of community and society in the remainder of the article.

Speeds and services

As shown in Figure 1, available network speeds tend to grow in jumps. The ARPANET used 56Kbps links for more than a decade. The Internet had 10Mbps Ethernet speeds commonly available from its inception in 1983, but used 56Kbps long distance links until about 1987, when T1 (1.544Mbps) started to be used. Since then, network speeds have begun to climb. A T3 (45Mbps) test network is in place, and faster wide area network speeds are expected. 100Mbps FDDI local area network speeds are now available. The speed increases shown for the years after 1990 are meant to be illustrative of a tendency for spurts every few years, with LAN speeds keeping somewhat ahead of WAN speeds. Such speed increases permit new services.

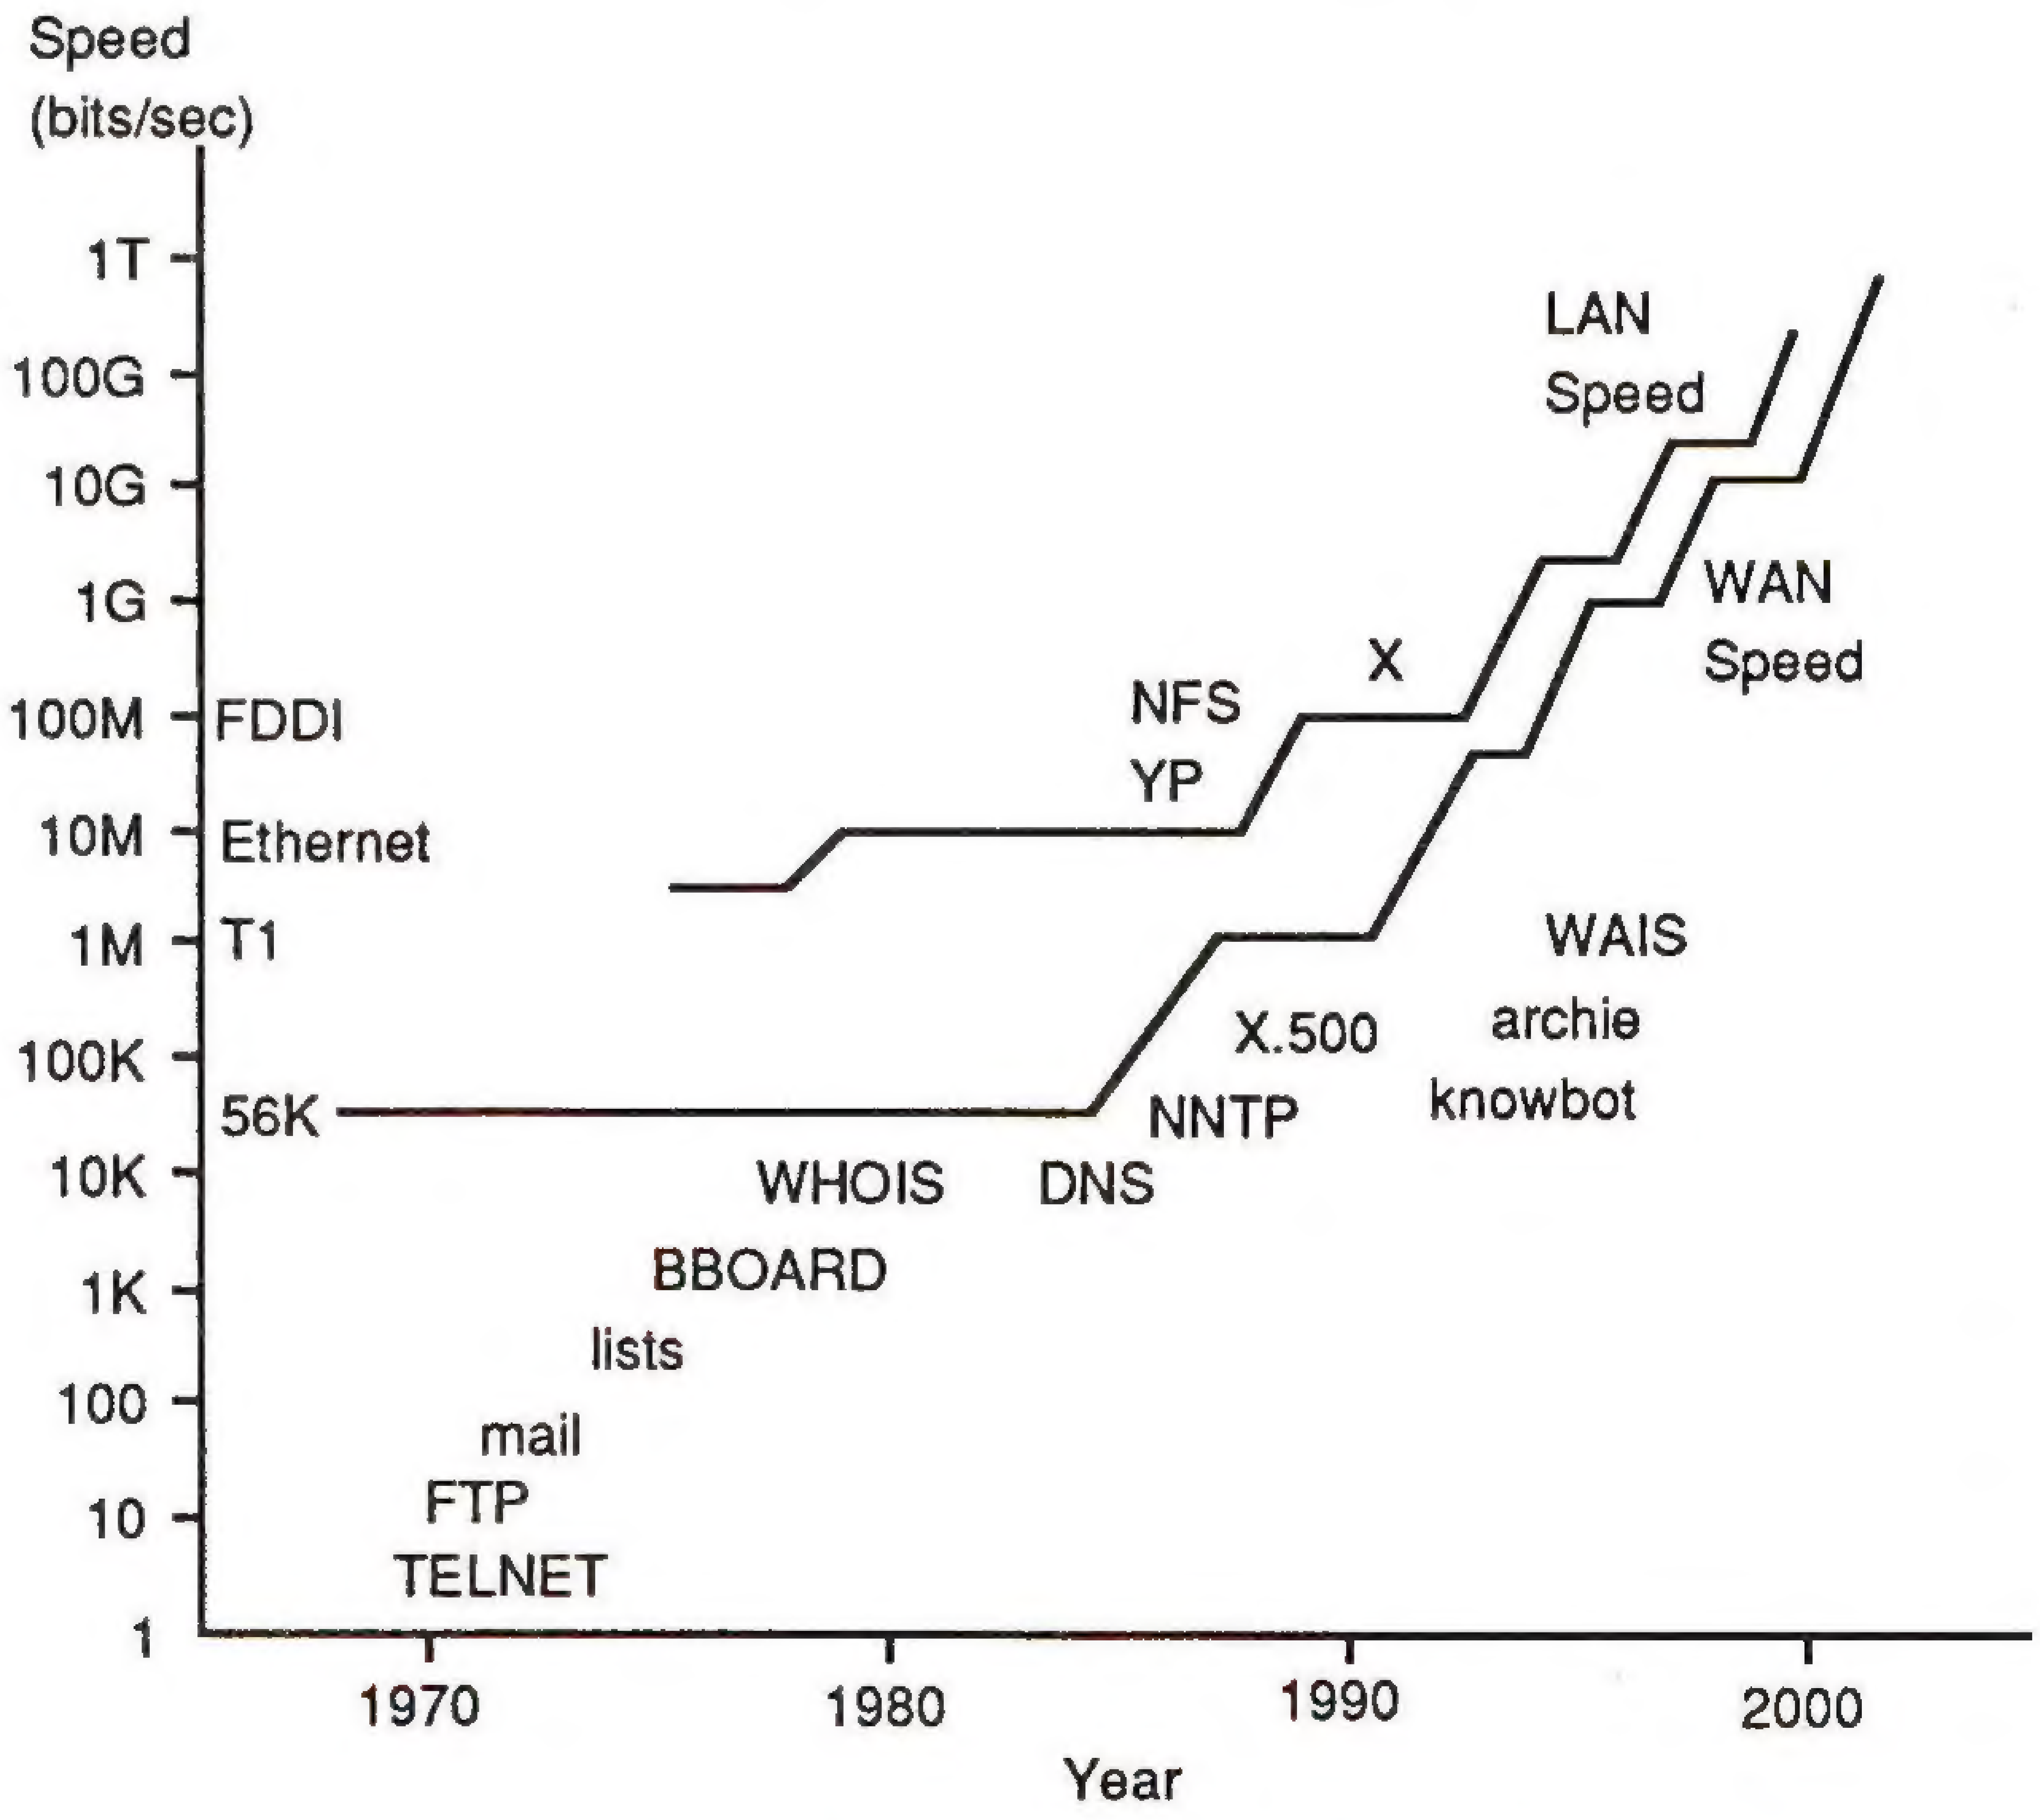

Figure 1: Network Speeds and Services

The ARPANET

The earliest multi-site packet switched network, the ARPANET, was intended for resource sharing. That is, the sponsoring agency, ARPA (now DARPA, the *Defense Research Projects Agency*), thought a network to connect its sponsored organizations would be less expensive than buying new large computers for each of them. The organizations could just use the network to log in on each others' computers and transfer files between them. [1]

This approach influenced the naming of network constituents: connected computers with users were called *hosts,* because people from elsewhere could log in to them as guests. It also influenced network protocol terminology, as processes or computers with resources to share were called *servers,* and processes or computers that used those resources were called *clients.* Services like FTP (File Transfer Protocol) and Telnet (remote login) were part of the original network plan, and were implemented soon after the first ARPANET nodes were in place in 1969. Host names were mapped to network addresses by a central file on each host, updated from a master copy at a centralized site.

Early ARPANET users quickly discovered they wanted to use the network to send messages to each other about the status of their projects. This idea of electronic mail was first implemented as an addition to FTP (perhaps more accurately described as a "bag on the side"). This seemed natural at the time, since mail was placed in a mailbox file per user, and the name of the mailbox file could be taken from the destination user's login name. Mail was soon the most-used service on the network.

Mailing lists

Users then discovered that they often wanted to send mail not just to specific people, but also to fixed groups of people, such as everybody participating in a particular implementation or planning task. These *electronic mailing lists* (or distribution lists) were implemented using aliases, that is, names that looked like mailbox names, but that were expanded by mail agents to lists of addresses for delivery. This illustrates that mail is something basically different from file transfer, since addresses in a mail alias may refer to users on any system reachable through the network, i.e., they are not limited to the sending or receiving (or controlling) host, as in FTP. That is, a *message transfer agent* (MTA) distinct from FTP is useful. When the ARPANET mail specifications were rewritten in the late 1970s, they were separated from the FTP specifications, and implementations of the new *Simple Mail Transfer Protocol* (SMTP) server were separate from the FTP server.

BBOARD

A next step was made when system administrators noticed that mailing lists involved a copy of the same message for each recipient on a host. This was a waste of disk space for large lists, since there were typically many users per host in those days. Many of those hosts were TOPS-20 or TENEX systems. (These ran on Digital DEC-20 or DEC-10 hardware; TENEX was developed by BBN from Digital's TOPS-10 operating system, and later revised by Digital as TOPS-20.) On such systems, a mechanism called BBOARD (for "Bulletin Board") became popular. Mailing lists could send one copy of a message to the BBOARD for each TOPS-20 or TENEX host. Users would then use the BBOARD command to select a BBOARD and read the messages in the mailing list it corresponded to.

Eventually, there were enough users using mail, lists, and BBOARDs that they wanted ways to find each other's mail addresses and other contact information. The *finger* and *WHOIS* services were invented for this purpose. Finger shows information about users on a single system. WHOIS is a centralized database for a whole network, with access methods.

All this was on the ARPANET, before 1980, with links running at 56Kbps. But this information was presented not only as a historical overview of a particular case.

From Technology to Community *(continued)*

The protocols and services of the ARPANET were direct ancestors of those of some other networks, especially the Internet. But there is a pattern here of resource sharing, mail, lists, groups, and user information services that recurs on many other networks, even unrelated ones. We will return to this pattern later in the article.

Local Area Networks

While the ARPANET was spreading all over the country and sprouting links to Hawaii and Norway, local area networks were being invented. Here we concentrate on *Ethernet.* The original Ethernet, as invented at Xerox had a theoretical maximum speed of 3Mbps, and was designed to throw away bandwidth. The later version from Xerox, Intel, and Digital, ran at 10Mbps, as did the protocol that was standardized as IEEE 802.3. Even though 10Mbps Ethernet was still designed to throw bandwidth away, 30% of 10Mbps is still 3Mbps, which is 53 times faster than 56Kbps. (It turns out that it really *is* possible to get 10Mbps transmission speeds out of 10Mbps Ethernet, but that is another story.)

56Kbps wasn't really fast enough (at least when multiplexed) to handle distributed file systems. Ethernet was. Xerox implemented a shared file system and a distributed name service, as did others vendors.

The Internet

Researchers involved with the ARPANET could see that one future of networking was interconnected sets of dissimilar networks, such as Ethernets connected by slower wide area networks of ARPANET-like technology. The *Internet Protocol* (IP) was invented to permit this, along with the *Transmission Control Protocol* (TCP), the *User Datagram Protocol* (UDP), and others in the TCP/IP protocol suite. In 1983, the ARPANET split into ARPANET (for network research) and MILNET (for operational use). Both ARPANET and MILNET became wide area backbone networks connecting local area networks into an internet, then called the ARPA Internet, now called just the *Internet.* All the old ARPANET services were available on the Internet as part of the new TCP/IP protocol suite.

The growth of the Internet was spurred by the release of the 4.2BSD (Berkeley Software Distribution) version of the UNIX operating system in 1983 and its revision as 4.3BSD in 1986. Meanwhile, new hardware technology allowed faster, smaller, and cheaper computers to spread.

NFS

New companies, such as Sun Microsystems, were formed to take advantage of these developments. Sun invented a *Network File System* (NFS), which allowed relatively transparent remote access to files, unlike FTP, where the users has to explicitly transfer a file before using local native programs with it. NFS was written to be used on top of UDP. It was made possible by the above developments, plus the availability of fast network technology such as Ethernet.

Such networked file systems brought a need for quick and distributed access across at least a local area network to information about users. For this purpose Sun provided YP (*Yellow Pages,* now known as NIS, for *Network Information Service*). NIS was designed for fast networks and is almost solely used on them.

Fast local area network speeds also permitted new variations on remote login. The *X Window System* was invented by MIT Project Athena around 1984. It, unlike NFS and NIS, is also fairly widely used over even fairly slow wide area networks, but its development was clearly spurred by fast network speeds. Similarly, the *Andrew File System* (AFS) can be used over slow networks, but was designed first on fast local area networks.

NSFNET

About 1984, proposals began to be drafted for a national super-computer access network, later called NSFNET and deployed in 1986. This became the main backbone network in the Internet. The NSF-NET backbone was implemented to use T1 (1.544Mbps) links about 1987. Experimental services such as packetized video and packetized voice began to be seen (some of these had been under development even on the old ARPANET, but weren't practical until higher speeds were available).

The ARPANET was retired from service in 1988 and 1989 because its link speeds were considered obsolete. Meanwhile, an experimental T3 (45Mbps) testbed network is already in place.

Population

Network speeds are not the only cause of invention of new protocols. Increasing numbers of networks, hosts, and users also have effects. Figure 2 gives very rough estimates of the user population of the ARPANET and then of the Internet from the beginning (1969) to the near future (2000). The ARPANET (pre-1983) figures are outright guesses. The Internet (1983–1990) figures were "computed" by multiplying the number of networks in the Internet by an average of 100 hosts per network and 10 users per host, or 1000 users per network.

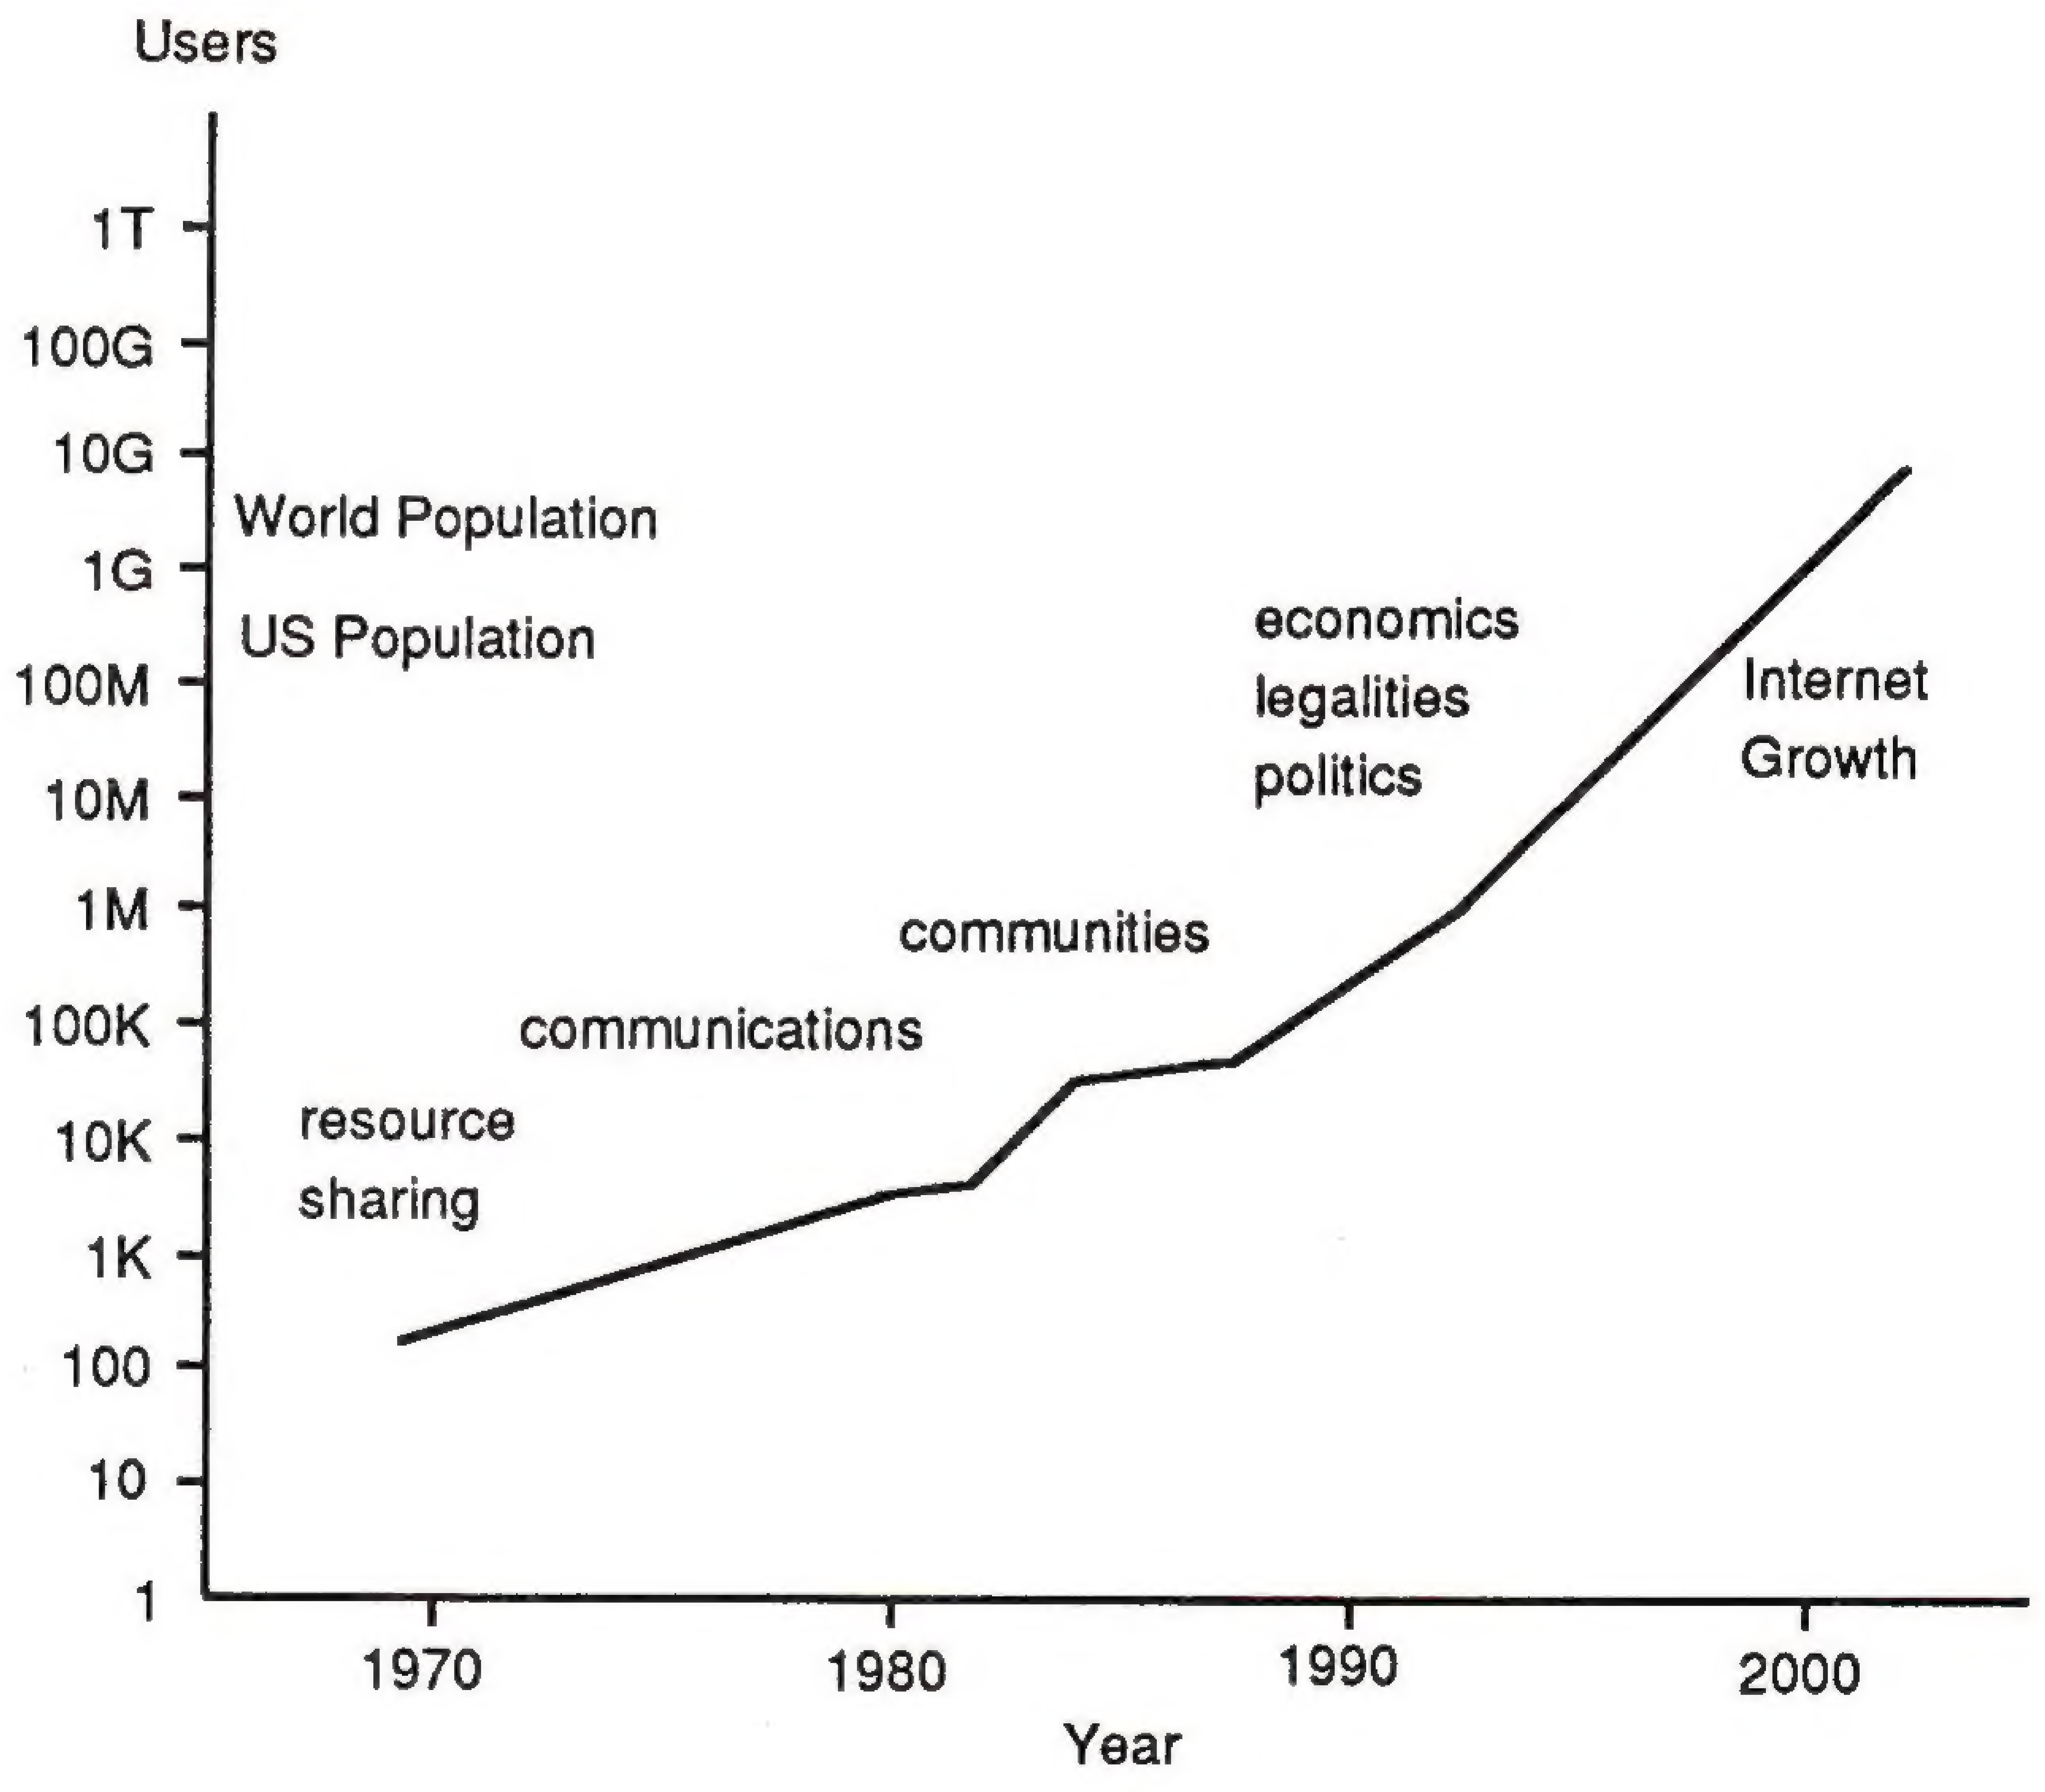

Figure 2: Internet Growth and Uses

The growth rate from about 1987 has been exponential, as the number of networks has doubled each year. The near future numbers (1991–2000) assume a simple continuation of this exponential growth. Obviously that can't continue forever, since we run out of people on the planet before 2000, but few who have tried to estimate Internet growth have shown any slacking in the exponential curve before the next five years or so. In fact, for the last year a doubling time of one year is clearly less than what is really happening.

continued on next page

From Technology to Community *(continued)*

Resource naming

Meanwhile, Internet protocol developers had to contend with not just one or two networks, but hundreds, and anticipated thousands in a few years, together with tens or hundreds of thousands of hosts. The old centralized tables for mapping hosts to network addresses would no longer be adequate. A distributed host naming service, DNS (the *Domain Name System*) was invented and deployed around 1984 to meet this need. This is one of the first examples of an Internet service that was clearly motivated by population pressures, not by higher available network speeds (e.g., DNS secondary servers just make a copy of the whole database). However, higher speeds later made the current very wide use of DNS practical. [2]

Similarly, the old-style WHOIS service eventually proved to be inadequate for large user populations. New services such as X.500 (the OSI Directory Service) have been implemented and deployed for this reason. No service has actually adequately met this need yet, and this is an active topic of network research and policy discussions. [3]

Increased numbers of networks and hosts have led to an alphabet soup of network routing and management protocols to deal with them, but those are peripheral to the main topic of this paper, because users do not see them directly; they merely support other protocols.

Resource access

In a network of tens of thousands of hosts, it can be very difficult to find things. For example, FTP supports an access method called "anonymous FTP," which allows anyone to connect to a host that supports it with FTP, log in as user *anonymous* with password *guest,* and retrieve files left there for general use. Source code for programs, binaries of programs, archives of mailing lists, protocol specifications, and a plethora of other information is available from anonymous FTP servers. But which anonymous FTP server host has the files you want? The traditional method of finding out is by polling or word of mouth; not very efficient.

A newer method is illustrated by the *archie* service of McGill University in Montreal. The archie server looks at a list of anonymous FTP servers, polls each one and retrieves an index from it, and keeps the indexes on a single host. Users may then connect to the archie server and examine these indexes to determine which anonymous FTP servers have the files they want.

Knowbots

Similarly, the *Knowbot Information Service* (KIS), or *Knowbot* for short, developed by the Corporation for the National Research Initiative (CNRI) automates searching servers that provide user directory information. A Knowbot can be configured to use the WHOIS service, the X.500 service, the finger service on an appropriate host, and other services, to fetch information about a user (or perhaps about a host, network, domain, or organization). Some formatting is done on the retrieved information to make it more legible, but no attempt is made to merge what comes from different servers, nor are relative values given.

Information access

It would be even more useful if computers could be used to automate choosing information, not just finding it. A step in this direction has been made with the *Wide Area Information Service* (WAIS), recently operational on the Internet. WAIS not only assists in locating servers, but also accepts rules from the user that help it determine what information to select. It can also be configured to keep looking and inform the user of new information.

Meanwhile, USENET *news*, which had developed independently on the UUCP dialup network at speeds of first 300bps, then 1200bps, 2400bps, and finally 10Kbps, had grown very popular by about 1988. The USENET news network had about 400 newsgroups at that time. Newsgroups are discussion fora somewhat similar to mailing lists, but kept per hosts like BBOARDs, and different from both in being much more widely distributed geographically (worldwide) and somewhat more independent of underlying hardware or software platform (most USENET hosts run UNIX, but there are also MS-DOS, VMS, CMS, etc. hosts).

The amount of traffic became the main problem in keeping the network running. Successively faster modem speeds had permitted it to continue to grow, but even 10Kbps became too slow. Fortunately, many of the main USENET hosts were also on the Internet, and the new NSFNET T1 backbone allowed even more traffic. Since USENET news over UUCP over TCP over IP was not particularly efficient or convenient, the *Network News Transfer Protocol* (NNTP) was invented to allow convenient distribution of news over the Internet.

People

There is clearly a pattern of networks permitting services (e.g., FTP) that are then used to build other services (e.g., mail). Faster network speeds then allow more transparent relatives of the earlier services (e.g., NFS) and also new services (e.g., X, although some say it is just a spiffy kind of Telnet). Population pressures combine with available speeds to permit and demand more transparent and distributed services (DNS, X.500, NNTP). Eventually, increasing population requires development of services to find and access other services (archie, Knowbot). Finally, services are needed that not only find and retrieve information, but that also actually interpret it for you (WAIS). But the more interesting aspect of this cumulative development of network services is what people do with them.

Resource sharing

The original goal of the ARPANET was resource sharing. This was also the goal that was used to justify funding of NSFNET, and is one of the goals being used to justify the proposed NREN (*National Research and Education Network*) [4].

Being able to explicitly access a computer located elsewhere is good and useful, as is being able to retrieve someone else's program or data. Resource sharing is essential to research and development or commerce. It is usually the first goal of R&D or enterprise networks. The first use of new network speeds is often more sophisticated resource sharing, e.g., NFS, YP, and X.

Even the USENET network developed from the UUCP mail network. UUCP stands for "UNIX to UNIX CoPy," and the underlying protocol does file transfers and remote job execution. Mail is implemented as a combination of the two. News was added later by the same kind of combination. But resource sharing preceded communication on UUCP and USENET, just as it did on the ARPANET.

Communication

People want to talk to people, not just machines. Computer networks rapidly become used for communication, thus known as *Computer Mediated Communication* (CMC). Mail was quickly invented on the ARPANET and UUCP. Even BITNET, whose underlying support mechanism emulates punch cards, has mail as its most widely used service [5].

Places

People want to talk to not just individual other people, but also to groups. Mailing lists develop on every network that has mail. People begin to depend on them as places to get information or hear interesting news.

continued on next page

From Technology to Community *(continued)*

Travel

Given places, people begin to travel between them. Given enough places, navigation is necessary. Sophisticated management services like BBOARD, LISTSERV (on BITNET), and USENET news (with its many interfaces) develop.

At this level of sophistication, the appropriate metaphor for use of computer networks may not be communications, with its familiar analogies of telephones, paper post, fax, radio, and television, but *travel,* with its immediacy of experience and its tendency towards total immersion. A later article may discuss rapture of the netways...

Communities

Once you can go to other places and come back, you begin to notice there are some places you feel more comfortable or get more work done. People begin to frequent these places, and some develop into *communities.* There is a sort of evolution from resource sharing through communication, places, and travel to community.

Computer networks have never been used solely for work. One of the earliest online communities was probably the SF-LOVERS mailing list, which was widely distributed over the ARPANET as early as 1978, despite never being sanctioned by any network authority, and several actual attempts to suppress it. There are hundreds or thousands of online communities today, many geographically distributed. These include not only the publicly advertised USENET newsgroups and Internet, BITNET, and UUCP mailing lists. Such communities can form whenever a group of people decide to start a mailing list.

Many networks have been justified on the basis of resource sharing, and many people say they use networks for communications. But some of those who pioneered networks such as NSFNET say that the real purpose was to form or facilitate communities. These goals are not necessarily contradictory.

These networked communities are different in some ways from other communities. By the nature of the services and networks that support most of them, they are distributed, asynchronous, and recorded. They are also diverse, and many members of them say they are egalitarian.

Some people worry that networked communities are "thin" communities, in that they do not involve direct human interactions as in "thick" communities such as a baseball team, a musical group, or a neighborhood church. Probably more networked communities tend to be thin than, for example, work groups in businesses. However, many networked communities lead to interactions among their members by other means, such as traditional media like telephone and paper post, and especially by travel for personal meetings. Perhaps it would be better to say that the networks *facilitate* communities.

Politics

Whenever you have communities of people, politics follows. The battles over the creation or charter of a USENET newsgroup can make old-time Chicago ward politics look tame. On a larger scale, the existence, funding, and access of the networks themselves have become political issues on local, national and international levels.

Politics in *networked communities* (or using network communications) may be somewhat different than traditional politics, even about networks. Traditional communication media tend to fall into two groups. Paper press, radio, and TV reach mass audiences, convey information, and leave impressions. Paper post, telephone, and fax reach individuals, are interactive, and can be used for actions. Computer mediated communications can do both. Could this lead to accountability of leaders? And perhaps empowerment of citizens?

However, I will note that while power may come from the barrel of a gun, as Chairman Mao said, it is often preserved by secrecy. In networking, secrecy is not power, and may not even be possible. Therefore, networking is subversive. That may or may not be the opposite of electronic democracy, but I will avoid that discussion here.

It is interesting to remember that this has all been made possible by an early grant from the Department of Defense.

Economics

One of the reasons networks have become politicized is that some of them, such as the NSFNET backbone, are partly government funded and thus influenced by government-defined acceptable use policies. Government funding is provided by taxpayers, who often have differences of opinion over what their tax dollars should go towards. One way out of that morass may be to privatize the networks, which would involve making them economically viable for commercial providers. This appears to be happening already.

Legal issues

Where there are differences over money or politics, we often find lawyers. And, in recent years, sometimes the Secret Service or the FBI. But those are other stories.

Summary

I hope I have sketched the bumpy slope up from bits to barristers. Networks may start as solely technological tools, but they don't stay that way if they survive. They develop places people go, which turn into communities, which develop politics, economics, and legal issues. The sum of all these things is a society.

Radio and television produced a different society. Computer networks will, too. Perhaps this time we can avoid a few mistakes.

Note: The numbers in Figure 2 for Internet population are purely speculative. The author hopes that readers of *ConneXions* will come forward with actual datapoints, especially for the early years.

[Ed.: This article is reprinted with permission from Matrix News, Volume 1, No. 2, May 1991.]

References

[1] Dern, D., "The ARPANET is Twenty," *ConneXions*, Volume 3, No. 10, October 1989.

[2] Mockapetris, P., "Introducing Domains," *ConneXions*, Volume 1, No. 6, October 1987.

[3] Benford, S., "Components of OSI: X.500 Directory Services," *ConneXions*, Volume 3, No. 6, June 1989.

[4] Roberts, M., "The NREN Takes Shape," *ConneXions* , Volume 5, No. 6, June 1991.

[5] Quarterman, J., "The Matrix—Computer Networks and Conferencing Systems Worldwide," Digital Press, 1990, ISBN 1-55559-033-5.

JOHN S. QUARTERMAN described the worldwide set of interconnected computer networks in the paper, "Notable Computer Networks," CACM, October 1986 (also in *Computers Under Attack,* Peter J. Denning, ed., Addison-Wesley, 1990), and in the book, *The Matrix,* Digital Press, Bedford, MA, 1990. He is a partner in Texas Internet Consulting and of Matrix Information and Directory Services, Inc., of Austin.

Current Status on Networking in Europe

by Bernhard Stockman, NORDUnet

Introduction

A lot of things are happening in Europe on the networking front. This article gives a brief overview of some of the activities. The current infrastructure has for a long time been dedicated to projects and organizations, national and international, which have resulted in a rather messy picture (see Figure 1). The situation is "suboptimal" to express it mildly. Several efforts are currently under way to create a more optimized arrangement, with a pan-European high-speed multiprotocol backbone in mind.

Figure 1: European connectivity, April 1991

Current de facto backbones

Different European networking organizations have come to bi- or multilateral agreements on line sharing. For example, the lines between Stockholm—Amsterdam and Amsterdam—CERN are shared by a group of national and international European networks forming a de facto backbone for involved European users. Similar agreements have been reached between other European organizations on other lines.

Connectivity to the US

The connectivity to the US is similarly "messy." A multitude of lines have been installed for projects and missions needs without having the overall connectivity in mind. Currently there is better connectivity between some areas *within* Europe *via* US networks which of course is an unacceptable situation. Ideally the transatlantic connectivity will be formed out of a few "fat pipes" connecting into key points in the US and Europe and shared by all networking organizations having interest in the transatlantic connectivity.

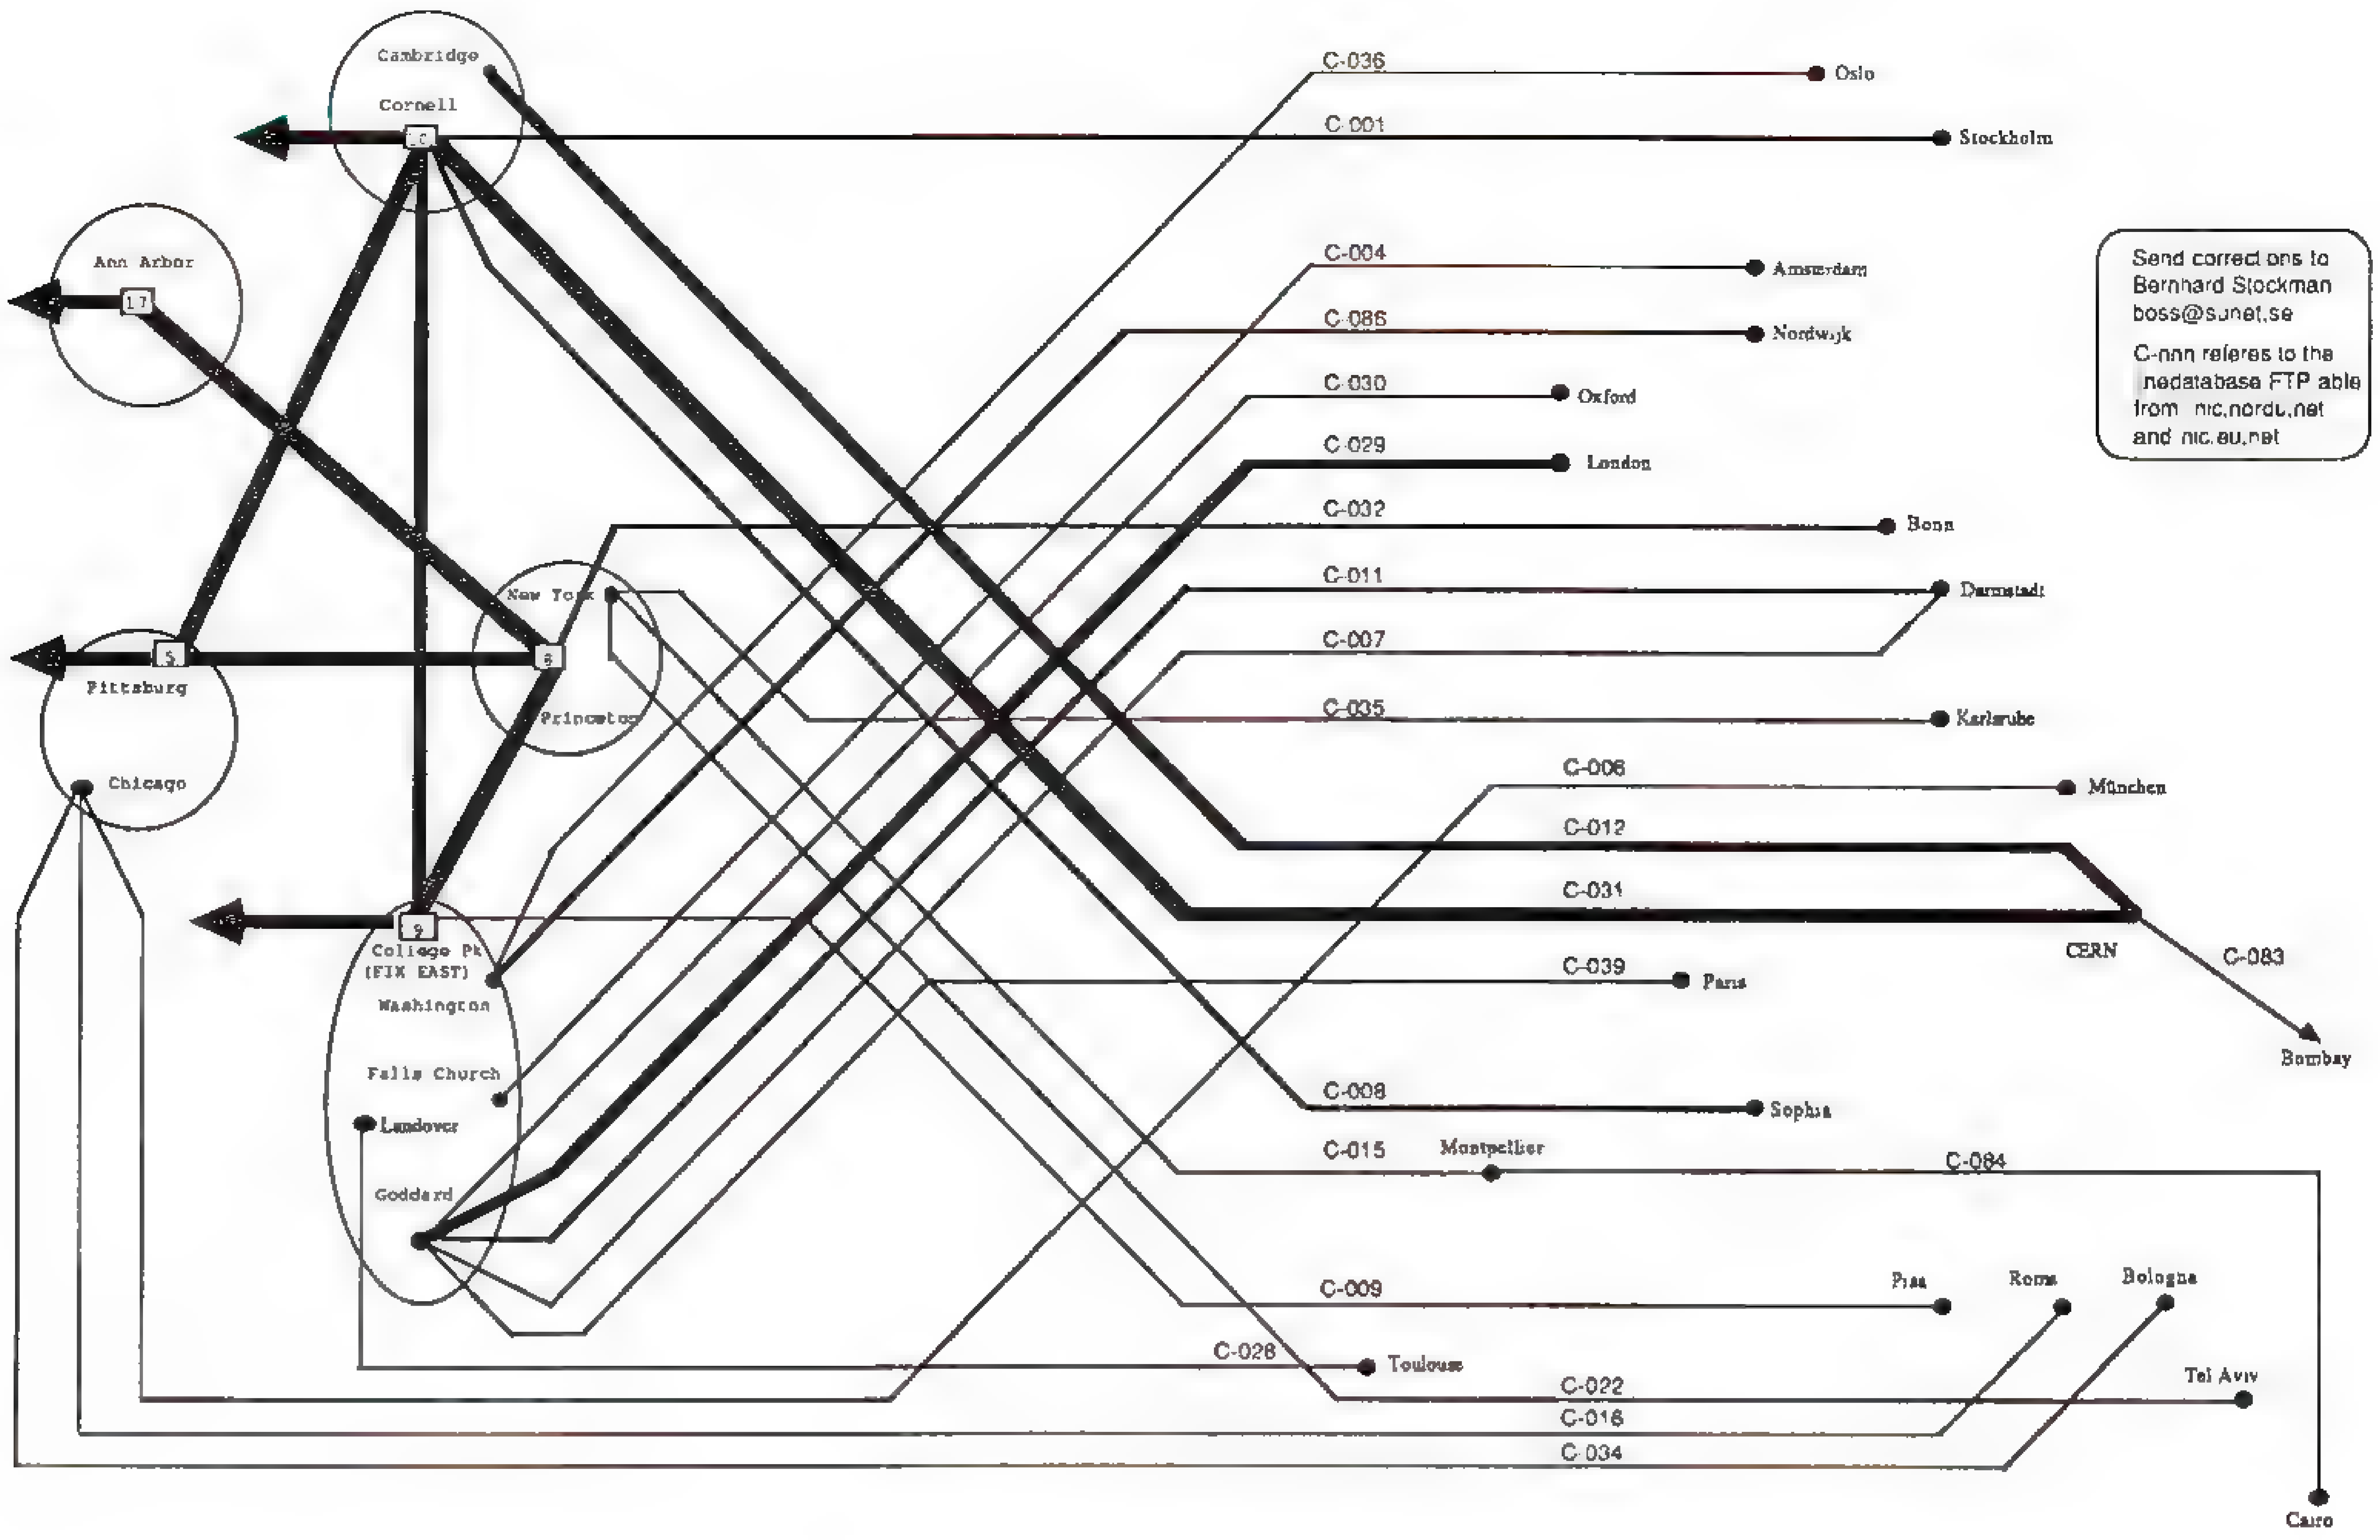

Figure 2: US–European connectivity, April 1991

IXI The European Commission has via the COSINE project initiated the *International X.25 Infrastructure* (IXI) project, which is a X.25-based backbone connecting European countries. The IXI backbone is currently set up with 64 Kbps connection points. An upgrade of the IXI backbone is being discussed. The IXI backbone is managed by the Dutch PTT Telecom BV in cooperation with other national telcos and appointed access point managers. Currently 19 European countries are interconnected via this backbone.

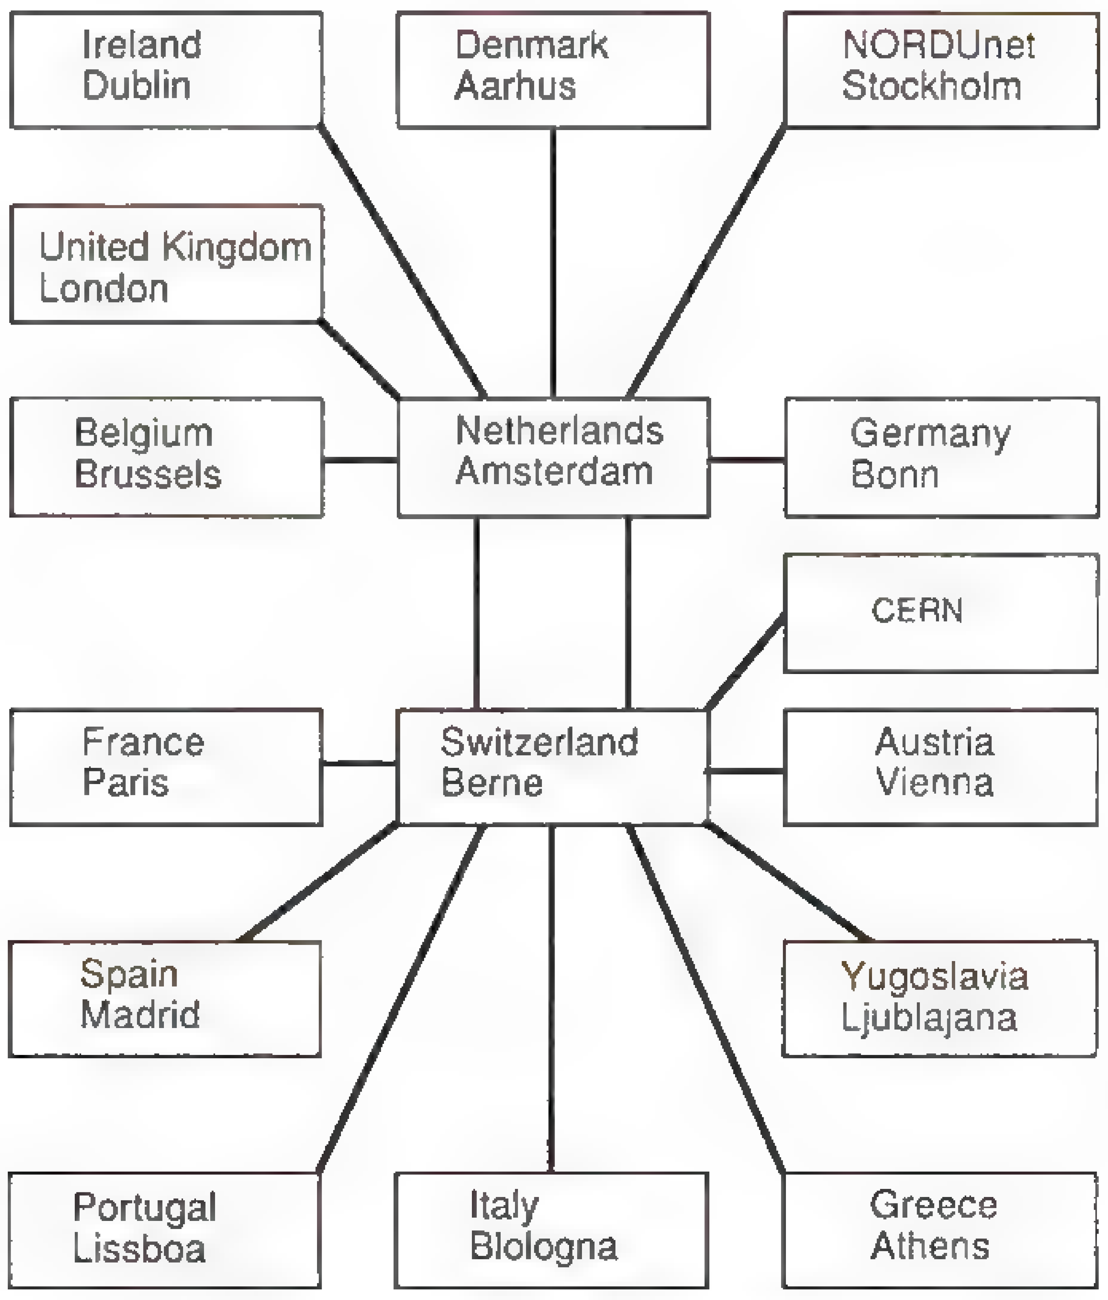

Figure 3: IXI backbone

EARN The *European Academic Research Network* (EARN), which is the European part of BITNET, has for long used its own private X.25 infrastructure within Europe. In cooperation with Digital Equipment Corporation an OSI experiment was performed using dedicated microVAXen donated by DEC and placed at key points in the EARN network. The outcome of this pilot was not very successful and the experiment has now been put to an end.

continued on next page

Networking in Europe *(continued)*

The EARN Operation Centre in Amsterdam responsible for the OSI experiment closed at the end of April 1991. Prior to this the private X.25 backbone was shut down and the line contracts terminated. EARN is now investigating the possibility of setting up a new backbone running NJE on top DoD-IP similar to the BITNET-II network, using the VMNET software from Princeton. EARN is cooperating with other organization in line sharing to connect the main European EARN sites. Regions not accessible via VMNET will be connected via NJE over X.25 on the IXI backbone described above.

RIPE

Réseaux IP Européens (RIPE) [Literally "Networks IP European"] is an organization that coordinates European IP activities. RIPE was formed in May 1988 in response to the pressing need for a total European connectivity plan with regards the DoD-IP protocol. Before this many ad hoc connections installed often resulted in unpredictable backdoor routing and other traffic problems. RIPE was accepted as a group within RARE during 1990 with the possibility of having some of its activities funded by RARE. Some examples of current work within RIPE includes:

RIPE NCC: The RIPE *Network Coordination Centre* (RIPE NCC), is an organization with the aim of putting some of the duties now being performed on voluntary basis into a formal structure with employed personnel. At the moment the structure of this organization is being defined with regards to its legal status, the possibility of having RARE (*Réseaux Associés pour la Recherche Européenne)* [Literally "Networks sharing for the Research of Europe"] as the funding organization, which duties the NCC should undertake, etc. One such task could be to act as an IP registry for Europe according to the procedure outlined in RFC 1174.

RIPE whois server: The RIPE *whois* database currently contains information on networks and persons. Four objects have been suggested for inclusion into the database: Autonomous Systems, Routers, Nameservers and Network lines. The database is accessible at `nic.eu.net` using the *whois* client software FTP-able from the same host in directory `ripe`. Other information such as logical *PostScript* maps on the European IP networks are also available via anonymous FTP from this host.

European part of the Domain Name System: The first DNS root server is now being installed in Europe at `nic.nordu.net` on NORDUnet. Tests are currently being performed for some of the zones. When these tests are successfully completed and the lines between NORDUnet and the US have been upgraded to 128Kbps, this first European root server will become fully operational.

EEPG

At the CCIRN and RARE meetings during 1989/90 the need was recognised for more concrete activities with respect to the planning and realisation of pan-European network services. During the Joint Networking Conference in Ireland in the spring 1990 the objective for the *European Engineering and Planning Group* (EEPG) were defined:

- Make a coarse estimation of near future networking needs.
- Make a coarse investigation of the current infrastructure and resources used.
- Propose some technical solutions for a pan-European high-speed multiprotocol backbone.
- Propose organizational solutions for managing this backbone.

Report

The EEPG has since delivered an interim report regarding the current infrastructure. Some items mentioned in the report include the establishment of an online database of lines within Europe and between Europe and other continents. The database can be accessed by anonymous FTP from `nic.nordu.net` or `nic.eu.net` in the directory `ripe/links`. The report also gives an estimation of the current resources used. For purely internal European international connectivity around $6 million is spend annually on telco charges. For intercontinental connectivity around $3 million is used, counting only the European half-channel. The cost for the IXI backbone is estimated at $4 million for 1991. All together around $13 million is spent annually on academic research networking in Europe. Two tables from the EEPG report give some ideas on European connectivity today:

Bandwidth (kbps)	Number of lines inside Europe	Number of lines outside Europe
2048.0	3	0
1544.0	0	1
1024.0	1	0
768.0	1	0
512.0	0	1
384.0	0	1
256.0	2	2
192.0	1	0
128.0	5	2
64.0/56.0	30	14
19.2	9	0
14.4	2	0
9.6	27	3
4.8	1	0

Table 1: Current frequency distribution of bandwidths for European international lines.

Country City	Total International Bandwidth	Total number of International lines
CERN	11516.8	28
FRANCE	2859.2	31
Montpellier	681.6	13
Paris	796.8	9
Lyon	1038.4	2
SWITZERLAND	5001.6	17
Berne	576.0	9
Lausanne	2112.0	2
Geneva	2057.6	2
GERMANY	2337.6	24
Darmstadt	880.0	10
Hamburg	806.4	4
NETHERLANDS	2099.2	30
Amsterdam	1372.8	19
Nordwijk	643.2	8
ITALY	2779.2	16
Bologna	2190.4	4
Rome	422.4	8
UNITED KINGDOM	1067.2	18
Oxford	241.6	8
London	585.6	3
SWEDEN	649.6	7
Stockholm	649.6	7

Table 2: Countries and cities with the major international connectivity

continued on next page

Networking in Europe *(continued)*

ECFRN

The *European Consultative Forum on Research Networking* (ECFRN) is a newly formed group from an initiative from the chairman of the ECC COSINE project. As organizations like RARE and COSINE only have national members, the idea was to form a group having representation not only from national networking organizations but also from international networks, The European Commission, national PTTs, etc. At the first meeting in Paris, March 8, around 75 people representing European users, governments and networks attended. Three main areas of interest were identified:

- A pan-European organization responsible for planning, implementing and running a pan-European network has to be found.
- The current funding of networking activities has to be more directed towards pan-European requirements in the future.
- Organizational and funding structures should be built to serve longer term planning. A 2Mbps pan-European backbone should be built immediately with possible near future upgrade to 34Mbps, with European gigabit testbeds in mind.

It was stated that a significant difference between US and European networking is the lack of continental carrier service providers in Europe and the relatively high telco tariffs there. These are big obstacles on the way to a pan-European high-speed network.

Eastern Europe

Many eastern countries are requesting IP connectivity to western Europe and the USA. The problem is that COCOM regulations still prohibit such connectivity to networks in the USA. In earlier requests from EARN it was stated from the US Department of Commerce that batch oriented traffic was approved while interactive traffic was not allowed. According to this NJE connectivity was installed from the University of Linz in Austria to Czechoslovakia, Hungary and Yugoslavia. Poland had already installed a 9.6kbps line from Warsaw to Copenhagen running NJE.

Poland has requested IP connectivity to DESY in Hamburg. Poland, Soviet Union, Hungary and Czechoslovakia have requested connectivity to HEPnet. EARN in Austria has received request from Czechoslovakia and Hungary to use the current EARN lines for IP traffic. Italy (Trieste) foresees IP connectivity to the same set of countries. Those countries have also expressed interest in participating in the RIPE initiative.

At the last CCIRN meeting these issues were discussed with US agencies present. The opinion was that the restrictions were unnecessary. The regulations may change, but for now no Eastern European traffic may enter the US part of Internet.

BERNHARD STOCKMAN graduated with a Master of Science degree in Electrical Engineering and Computer Systems from the Royal Institute of Technology in Stockholm, Sweden in 1986. After a couple of years as a researcher in distributed computer systems he was employed by the NORDUNET and SUNET Network Operation Centre where he is responsible for network monitoring and traffic measurement. He is participating in international cooperation as chairman of the IETF working group on Operational Statistics and also in similar groups in RIPE, RARE and NETF. For the past year he has been working in the European Engineering and Planning Group (EEPG).

Another use of the Internet: Libraries Online Catalogs

by Billy Barron, University of North Texas

Introduction

Over the last several years, library online catalog systems have started becoming available over the Internet. At the latest count, 192 libraries are currently accessible over the Internet. This number is growing at the rate of more than one new system per week. This article will look at the current and future use of the Internet library computer systems and the information available on them.

Uses of the Internet Libraries

Many different groups of people use the Internet libraries including librarians, faculty, students, and researchers. Some of the uses of the Internet libraries are:

- The ability to search other libraries for books that you wish to acquire through *Inter Library Loan* (ILL). Also, this library access allows you to check the status of the book before submitting an ILL request [3].
- The ability to check holdings of other libraries before going to research there. In addition, the location of the book or other materials can usually be determined in advance. Finally, the trip is not wasted if the library computer goes down while the researcher is there [1].
- Some ILL librarians use the library systems because locally kept information is usually more accurate than the national databases, such as OCLC and RLIN [7]. However, other librarians point out that though the Internet libraries are free and services, such as OCLC, cost money, the use of the Internet libraries require greatly more staff time than OCLC does, and the staff time costs more than the OCLC charges would [6].
- The library systems allow researchers to work at almost any time of the day instead of just the times when the library is open.
- Colleages at different sites working on collaborative projects can check each others libraries' holdings.
- Librarians can access RLIN and OCLC's EPIC directly over the Internet.
- Research and faculty members who are changing universities can look over their new university's holding in advance.
- People who do not have access to OCLC or RLIN can still find books for Inter Library Loan.
- The system can be used for verifying information for library acquisitions when some information is questionable [8].
- Library catalogers can review how other libraries have listed materials for consistency [8].
- Data Research Associates has made cataloging information available via FTP to `DRA.COM`.

Current state of the Internet Libraries

Presently, most major and many smaller universities have their libraries available on the Internet. More and more are coming on-line all the time. Although this is an encouraging sign, the growth does present some problems.

continued on next page

Libraries Online Catalogs *(continued)*

Unfortunately, there are many different library software packages in use and all of them have slightly different user interfaces. Computer knowledgeable users can handle the different interfaces pretty easily. However, for others, this causes problems and results in less usage and wasted time.

Users on IBM mainframes running IBM's TCP/IP package, *FAL,* are unable to access library systems that require a VT100 terminal. Fortunately, very few systems require VT100 though many have it as an option. Some third-party products exist that will add the VT100 support, but they are "fairly hokey" [4].

TN3270

Users on some machines (such as Sun) do not have TN3270 [12, 13]. Therefore, they can not use the libraries available at many IBM mainframe sites. In many cases, this is a solvable problem. For UNIX systems without TN3270, a TN3270 package is available for anonymous FTP on `ucbarpa.berkeley.edu` in directory `/pub/tn3270`.

Milking machine

Another solution for the IBM mainframe-based library requiring TN3270 can be implemented on the IBM side. This solution is known as a *milking machine.* Many terminal servers have a feature known as *Reverse Telnet.* The serial lines of the terminal server can be connected to a protocol converter such as a 7171. Then the VT100 user can Telnet through the protocol converter which will translate between VT100 and 3270 emulations.

The milking machine has another important use in this arena. Some of the library computer systems do not have TCP/IP support. This milking machine can again be used. Instead of hooking up the terminal server lines to a protocol converter, it can be hooked into a serial port on the computer to provide Telnet access into the library system.

The major current limitation is the lack of a way for the library systems themselves to communicate with each other. Unfortunately, no easy solution to this problems exists today.

Allowing access via the Internet could possibly overload the library computer to the point where it is unable to service its local users. Fortunately, it is possible for the administrator of the library computer to limit the number of simultaneous Internet users [9]. The policy and political objectives of connecting the library computer to the Internet are usually greater and should be considered before the technical issues [5].

Future of the Internet Libraries

While it is obvious that more and more libraries will be on the Internet, many other changes will be taking place also. Any major improvements, however, involving the Internet will require some new protocols.

Z39.50

The *National Information Standards Organization* published the *Z39.50 Information Retrieval Service Definition and Protocol Specifications for Library Applications* in 1988 [2]. Z39.50 is a protocol that searches and retrieves information from a remote computer system. While it is designed for bibliographic data, it is general enough to support many different types of information. Many groups are working on the implementation of Z39.50 both on TCP/IP and OSI. The benefits of Z39.50 might include remote queries, transmission of cataloging information, and Inter Library Loan requests over the network.

Another likely change in the future of library computer systems is the switch from a text based user interface to a graphical user interface based on The X Window System. Also, it is likely that Hypertext features would be included in such an interface. Though these new user interfaces have been discussed, I am not aware of any actual implementations existing today. Successful implementations may require that the Z39.50 protocol comes into common use first.

To take advantage of these future developments, librarians need to become more network knowledgeable. Likewise, the network implementors need to become more aware of library computing issues. The common goal should be to make the library catalogs more accessible and easier to use.

Information resources

In the Internet library information community, there is a definite split into two camps on some issues. The first group believes that all information should be approved with the institution before publication. The St. George Guide, Internet Resource Guide, and LIBTEL follow this model. The other group believes that all information should be printed unless requested otherwise by the institution. Of course, the information is verified for accuracy. The main belief is that if a site does not want Internet users accessing its system, the site should employ security measures such as routers to limit access, instead of depending on security by ignorance. Also, it is felt that the verification process slows down the spread of information. The Barron guide, HYTELNET, and CATALIST follow this model. In any case, both groups do make important contributions to the information resources that are available on the Internet libraries.

Documents

Several useful documents exits to aid access to the Internet libraries, and they are listed below. I attempt to keep a complete archive of these documents on `VAXB.ACS.UNT.EDU` accessible by anonymous FTP.

UNT's Accessing On-line Bibliography Databases by Billy Barron, University of North Texas. This is a clear and concise guide to library systems available on the Internet, JANET, and THEnet (Texas Higher Education network) [14]. Available via anonymous FTP on `VAXB.ACS.UNT.EDU` in the `library` directory. Submissions for new or updated entries may be sent to `billy@vaxb.acs.unt.edu`.

Internet—Accessible Library Catalogs & Databases by Art St. George, University of New Mexico, and Ron Larsen, University of Maryland. A guide to Internet and JANET library systems and campus-wide information systems. Available via anonymous FTP from `ARIEL.UNM.EDU` in the `library` directory. Submissions for new or updated entries can be sent to `stgeorge@unmb.bitnet`.

Internet Resources Guide by the NSF Network Service Center (NNSC). Chapter 2 contains detailed information on a limited number of Internet library systems. This information is also summarized in the above two listed guides. Available for anonymous FTP from `NNSC.NSF.NET` in the `resource-guide` directory [10].

AARNET Resources Guide. This is the equivalent to the Internet Resources Guide for AARNET, the Australian part of the Internet. This information is also summarized in the Barron and St. George guides. Available for anonymous FTP from `AARNET.EDU.AU` in the `pub/resource-guide` directory [11].

continued on next page

Libraries Online Catalogs *(continued)*

OPACS in the UK: A list of interactive library catalogues on JANET compiled for the JANET User Group for Libraries by the University of Sussex Library. This is the original source for accessing JANET libraries. It is included verbatim in the St. George guide and summarized in a format consistent with the rest of the information in the Barron guide.

Programs

There currently exist two types of computer programs to aid users in accessing Internet libraries. The first presents connection and usage information in a easy-to-use Hypertext format. The other type actually connects you to the Library system automatically.

HYTELNET by Peter Scott, University of Saskatchewan. HYTELNET is an MS-DOS TSR that presents connection information on the Internet and JANET libraries and Campus-wide Information Systems in a hypertext format. Available via anonymous FTP in the `library` directory on `VAXB.ACS.UNT.EDU`. Suggestions for improvement should be sent to the `LIB_HYTELNET` list discussed below.

CATALIST by Richard Duggan, University of Delaware. CATALIST is an MS-Windows program that presents the Barron guide in a hypertext format. In the future, this product will automatically connect you to the library systems also. Should be available via anonymous FTP in the `library` directory on `VAXB.ACS.UNT.EDU` by the time this article is printed. Suggestions for CATALIST can be sent to `duggan@brahms.udel.edu`.

LIBTEL by Dan Mahone, University of New Mexico. LIBTEL is a program which will automatically connect you to the Internet libraries listed in the St. George guide. UNIX and VMS versions are available for anonymous FTP in the `library` directory on host `VAXB.ACS.UNT.EDU`. Suggestions for this program should be sent to `dmahone@hal.unm.edu`. This program can be seen in action via Telnet to `bbs.oit.unc.edu` and choosing item 9 on the main menu after you register for the BBS.

Mailing lists

While at least a dozen mailing lists exist on library topics, two of them stand out as essential forums for discussing library systems on the Internet:

PACS-L moderated by Charles Bailey, University of Houston. The Public Access Computer Systems List (PACS-L) covers all aspects of patron computer use in libraries. This is by far, the premier library computing mailing list. Major updates to the documents and programs listed above are normally announced on this mailing list. To subscribe send mail to `LISTSERV@UHUPVM1.BITNET` with a body of `"SUBSCRIBE PACS-L firstname lastname"`. Posts are mailed to `PACS-L@UHUPVM1.BITNET`.

LIB_HYTELNET moderated by Peter Scott, University of Saskatchewan. LIB_HYTELNET's primary purpose is to discuss the HYTELNET package. However, the majority of the traffic is about new library and other information systems that are available on the Internet. To subscribe, send a message to `SCOTT@SKLIB.USASK.CA`. Messages for the list are sent to `LIB_HYTELNET@SASK.USASK.CA`.

Conclusion

As we have seen, the Internet library systems provide many useful services to users and the methods of access are well documented. At the current time, these systems have some major limitations.

Fortunately, plans for overcoming these limitations are in the works and should be implemented over the next several years. Hopefully, by the turn of the century, the Internet libraries will be an extremely easy to use yet powerful service.

Acknowledgments

Special thanks to Peter Scott for editing assistance and some useful information. Thanks go to Philip Baczewski, Nancy Gusack, David Lippke, and Clifford Lynch for providing some much needed information and assistance that made this article possible.

References

[1] Ausdel, Karl, "RE: Internet Access to OPAC," Electronic Mail Posting, Public Access Computing Systems Forum, (`PACS-L@UHUPVM1.BITNET`), 11 April 1990.

[2] Hinnebusch, Mark, "Z39.50: Where is it and Who Cares?" *The Public Access Computer Systems Review,* Volume 1, No. 2, 1990, p. 62–66.

[3] Kobulnicky, Paul, "RE: Internet Access to OPAC," Electronic Mail Posting, Public Access Computing Systems Forum, (`PACS-L@UHUPVM1.BITNET`), 11 April 1990.

[4] Lippke, David, "Re: Telnet on VM," Private Electronic Mail, 11 May 1991.

[5] Lynch, Clifford, "Linking Library Automation Systems in the Internet: Functional Requirements, Planning, and Policy Issues," *Library Hi Tech,* Volume 7, No. 4, 1989.

[6] Maehl, Darren, "Re: Internet Access," Electronic Mail Posting, (`PACS-L@UHUPVM1.BITNET`), 20 July 1989.

[7] Perry, Andy, "National Bibliography Databases," Electronic Mail Posting, (`PACS-L@UHUPVM1.BITNET`), 20 July 1989.

[8] Scott, Peter, "Re: Reasons to use Internet OPACS," Private Electronic Mail, 13 May 1991.

[9] Tennant, Roy, "Still More on Internet OPACS," Electronic Mail Posting, (`PACS-L@UHUPVM1.BITNET`), 3 August 1989.

[10] Roubicek, Karen, "Information 'Products' from NNSC," *ConneXions,* Volume 5, No. 6, June 1991.

[11] Huston, Geoff, "Profile: AARNet," *ConneXions,* Volume 4, No. 3, March 1990.

[12] Minshall, Greg, "The 3270 Environment," *ConneXions,* Volume 1, No. 6, October 1987.

[13] Minshall, Greg, "3270 over Telnet," *ConneXions,* Volume 1, No. 7, November 1987.

[14] LaQuey-Parker, Tracy, "Profile: THEnet," to appear in *ConneXions,* 1991.

BILLY BARRON received his B.S. (1986) from North Texas State University and M.S. (1988) from University of North Texas. For the past five years he has worked for the University of North Texas in various capacities and now is the VAX/UNIX Systems Manager running VAX and various UNIX computer systems. In addition to typical system management duties, he is heavily involved in networking for the campus and has compiled one of the major guides to library computer systems on the Internet. He can be reached as `billy@vaxb.acs.unt.edu`.

Commercial IP providers establish CIX gateway

by Daniel P. Dern

Commercial, company-to-company internetworking using TCP/IP got a major boost on March 25, 1991, with the announcement of the first *Commercial Internet Exchange* (CIX), linking the three leading PDIs (*Public Data Internets:* term and acronym courtesy of your friendly Internet correspondent—other contenders include "Public Protocol Networks" and "IP Commercials"). The CIX (pronounced "kicks"), a jointly-operated resource, became operational June 1, 1991. Located in San Francisco, it consists of high-speed routers from Cisco Systems, Inc. (Menlo Park, CA).

Participants

The San Francisco CIX will initially interconnect *CERFnet,* run by General Atomics (San Diego, CA); *PSInet,* from Performance Systems International Inc. (Reston, VA); and *AlterNet,* from UUNET Technologies Inc. (Falls Church, VA). All three networks currently operate T1-speed backbones. AlterNet and PSI are nation-wide in scope, with some international connections; CERFnet serves a large portion of California.

Previously, inter-PDI connections either had to be made via the National Science Foundation's *NSFNET* backbone, or not at all. NSFnet usage is restricted to research and education-related traffic; while some inter-PDI traffic may qualify, other is clearly commercial in nature—and many corporate networkers are legitimately leery of relying on government-funded networks as part of a business resource. So the CIX avoids the "use/don't use" NSFNET question—and joins PDI communities and resources into a larger common pool. Peer-to-peer network connections are also springing up, e.g., NEARnet-to-Alternet, increasing non-NSFNET-dependent connectivity still further.

Network managers in organizations currently using these commercial TCP/IP-OSI networks, which includes Hewlett-Packard, Unocal, Trusted Information Systems and hundreds of others, all agree that that the CIX enhances the already-signficiant value of public protocol-service-oriented networks.

According to Peter Ho, Unocal Corporation (La Brea, CA), a billion-dollar-ish-sized energy company serving the West Coast and mid-west U.S., "We're currently using CERFnet to let the seismic engineers in our regional sites connect the Sun workstations on their desks to our supercomputing facilities, to run their simulations. To do this you need a low-latency, high-bandwidth connection."

The CIX will make it easier to connect places beyond CERFnet, suggests Ho. But more important: "We'll be able to communicate directly with our vendors. We can get microcode updates directly FTP'ed [copied via TCP/IP's File Transfer Protocol] to our systems, instead of waiting for a tape shipment. And we can exchange e-mail with organizations on the other networks, without worrying about violating 'acceptable usage' policies of some connecting government backbone."

A win

"For TIS, the CIX is an immediate 'win,'" states Steve Crocker, VP at Trusted Information Systems (Glenwood, MD). "Our Mountain View, California and our Glenwood, Maryland offices are on AlterNet. Our Los Angeles office connects to CERFnet, through the Los Nettos regional network. The CIX will create direct connectivity between the Los Angeles office and the other two, without involving the NSFNET backbone and all the potential traffic restrictions that involves."

"We don't generate a lot of traffic among our three offices, but when we do, we need T1 speed and responsiveness to connect the workstations at our three offices for wide-area filesharing, remote login under X-Windows, and other applications," Crocker says. "You can do this on slower facilities, but not well. For about the price of a 9.6Kbps leased line, connections to AlterNet give us the speed and latency we need."

Services

All three networks sell direct IP connections to their topologies at speeds from 9.6kbps to T1, representing nearly 100% of commercial TCP/IP and OSI services currently sold in the U.S. Other connection options (which vary by vendor and location) include dial-in terminal ports and LAN interconnect; PSInet also supports DECnet, Novell IPX/SPX and AppleTalk (in selected locations).

Cost

The protocol-oriented connectivity sold by the networks, typically priced as a fixed-rate based on the speed of each connection, is seen as offering the high bandwidth, low latency and reliability (and turnkey, managed service) at a controlled, affordable price, versus alternatives such as X.25, Frame Relay, leased lines or dial-up. For example, AlterNet and PSInet charge roughly $1,000 and $1,800 per month for a 56Kbps connection to their respective networks. This covers unlimited, bi-directional traffic through the DDS circuit as well as transiting to other networks.

The service price can include on-site router hardware, and help from network engineers to order the leased lines between customer sites and the PDN's nearest Point-of-Presence (POP), plus installation and integration.

Applications

The CIX will let customers of these networks access the other networks directly, for applications such as:

- Cooperative program development
- Customer technical support
- Participation/awareness of technical standards
- Access to online news and information
- Professional and staff development through immediate communications with leaders in the field through e-mail, news groups, and transferring late-breaking information and technology
- Inter-office communications and product briefings
- Working with vendors, suppliers and customers electronically
- E-mail and file transfer for price lists, product information, bug fixes, code patches, etc.
- LAN-class file-service under NFS
- Document conferencing
- Remote login for supercomputer access
- New services like "pay for view" databases
- Remotely-handled system administration and maintenance, software porting, etc.
- File transfer using TCP/IP FTP
- Electronic mail
- Opening concurrent windows to hosts in multiple offices

continued on next page

IP providers establish CIX gateway *(continued)*

Acceptable Usage

Prior to the CIX, traffic seeking to transit from one commercial network to another, e.g., AlterNet to PSInet, might have had to flow over the NSFNET backbone—which, depending on the nature of the traffic, might violate the NSF's "acceptable usage" policies against commercial traffic. Or, equally problematic, connections involving the NSFnet backbone might be refused.

Individual regional networks, often self-funded, could often accept commercial traffic within its boundaries without regard to its content or intent. But traffic headed to other regionals via the NSFnet needed to fit within NSF's usage guidelines. Regional guidelines—and interpretations—for transit traffic vary greatly, which doesn't help matters, either.

The CIX removes this selective, often unpredictable and inconsistent non-connectivity, and the barrier to greater business-oriented internetworking which it has posed for current users. For as-yet-unlinked organizations, the CIX simplifies the "which network" question, and makes company-to-company networking more tempting than ever.

Commercial customers have been increasingly baffled and frustrated by usage restrictions for the NSFnet backbone and many regional networks. "The Fortune 1000s, financial service companies, insurance and the like are used to being able to buy connectivity that doesn't have restrictions, whether for phone, data, or courier services," points out John Eldredge, Director of Sales and Marketing at PSI.

Commercial opportunities

"The CIX opens up a lot of commercial opportunities," says Dr. Vinton Cerf, a long-time contributor to the evolution of the Internet. "It will be a lot easier for organizations to interact with each other."

PSI and UUNET's initial solution, for their own customer bases, was to build nationwide backbones, giving their commercial customers nationwide connectivity without involving the NSFnet or other restricted backbones.

With more than one commercial IP network, however, plus wider customers bases on regional networks such as NEARnet, the "can/can't" transit question reopened. Many networks had already begun one-on-one connections, e.g., NEARnet and AlterNet. Now, the commercial providers have put together a generalized solution, which can also accommodate further members—the CIX.

Seamless internet

For CERFnet, AlterNet and PSI customers, the three networks will appear as one large, seamless internetwork. User systems on one network will be able to establish internetwork connections to systems on the other two. The greatest benefit, all agree, of the CIX, is increasing access to the people, information, and resources on each individual network.

"The CIX is very significant," says John Gong, Consulting Manager for Corporate Telecommunications at Hewlett-Packard. "We probably would have connected much sooner if we'd known it was coming, because we perceived the separateness of the commercial vendors as a roadblock. [HP joined PSInet in April 1991, connecting its headquarters at T1.] The CIX makes connectivity far more valuable—and eliminates the 'join two, or pick one' decision."

According to Susan Estrada, Executive Director of CERFnet, "We know that our networks are only as valuable as the people they connect. The CIX is part of that. We're three competing networks, yet we're working together to do things for our customers."

Steve Wolff, Director of Networking at the National Science Foundation, agrees. "When I was growing up, in Pennsylvania, we had two systems: Keystone Telephone Company and Bell of Pennsylvania. They didn't interconnect—so you had to have a phone and phone number from each, to be sure you could reach everyone. Similarly, without a CIX, you have only separate user communities, one for each commercial Internet carrier, and that diminishes the value of each to all their subscribers. The value grows as a power set—the total number of subsets—of the user populations."

Commercialization of the Internet, and development of commercial Internet-class service, has been a major source of discussion within the Internet community in the past year, fueled strongly by the growing use of UNIX and TCP/IP by commercial and corporate users.

A commodity

"The CIX represents another step in the direction which NSF has been pushing," says Wolff. "We want see networking become a commodity that *anyone* can procure, without a lot of special conditions or a staff of wizards."

DANIEL P. DERN is a frequent *ConneXions* contributor. Dern is a free-lance technology and business writer based in Watertown, Mass., who has written about Internet topics from OSPF/IS-IS routing to John Romkey's Internet Toaster. He can be reached as `ddern@world.std.com`.

INTEROP is coming soon!

INTEROP® 91®
7–11 October 1991 • San Jose, CA

By now you should have received the *INTEROP 91 Fall Advance Program.* If not call 1-800-INTEROP or 415-941-3399. The conference and exhibition will take place October 7–11 in San Jose, California. Choose from 45 conference session, 29 tutorials, lots of Birds of A Feather (BOF) sessions, and see the industry's premier networking exhibition where over 150 vendors demonstrate the latest and greatest computer and communications technologies on the show network.

San Jose Convention & Cultural Center, site of INTEROP 91 Fall.

Internet Resource Accounting

by Don Hirsh, Meridian Technology Corporation

Background

A wide variety of pressures are bringing long-term monitoring—and with it the ability to measure policy compliance and resource usage—to a higher priority in the scheme of network operations technology. The most obvious of these pressures is the changing nature of network management, from interesting engineering problem to rationalized utility administration. Several formal groups are at work with contributions ranging from economic policy to architecture and implementation. Interested parties abound, diverse opinions are not in short supply. The topic is both important and timely. Policy has preceded implementation with the consequence of many words and few programs. But where there is smoke there is fire, and the pace of experimentation is increasing. Changes in day-to-day policy are still a long-way off.

Networks aren't free

You'll have to forgive me for stating the obvious: over the past several years we have collectively come to understand that computer networks of any size are *not* free. We all know that—we know it intellectually...of course it's obvious. But enterprise-wide networks and internets and campus-networks have this unpleasant cost factor associated with them. I do not mean to be patronizing towards the readership of this journal, this is one of those points that is so obvious one doesn't even think about it.

How can we put this delicately? Those of us in networking at large institutions with broad goals have worked at developing exciting, enabling technology. Cost has been a secondary issue. No more. At all levels of the Internet, changing patterns in funding, an increasingly operational charter for support and maintenance organizations, escalating demands for quality service—all are serving to focus attention on the budget. More importantly for the future of the Internet will be the questions of who will pay for what services and how. For years and years we have justified the pleasures of network research by pointing out the providential utility of the Internet and our various sponsors have taken us at our word. With this reasonably mature technology comes the operational accountability that goes with this practical research discipline.

A utility

Networks of every sort are now viewed as an *information utility*—for the knowledge worker networks are as ubiquitous as the telephone or electricity. As they are seen to be utilities, the need to elicit efficient behavior from network consumers and providers surfaces. Hence we need capacity planning, resource allocation, cost recovery to allow for both current operations and future provision, and feedback that allows network consumers some window onto their utilization of a finite resource.

It's not that we're going to propose solutions for any of these complex problems. Rather, the problems associated with a loose confederation of activities changing to an organized utility focuses a debate upon useful policy. The conception and implementation of these policies is influenced by the available tools (remember the adage that goes "if you have a hammer, everything looks like a nail"). The current set of network management tools are inadequate for the task. So while a broad, important debate on Internet policy gains momentum and volume, an *Internet Engineering Task Force* (IETF) working group is specifying a set of interim standards for *Internet Accounting.*

While the focus of the IETF is on the Internet itself, the same set of operational pressures are coming to bear on network managers and administrators from every size and type of organization.

Precedents from daily life

There are several precedents that will spring into most people's minds in looking for models of internet resource allocation. Though suggestive, internets are more complex than any of the utilities that we will pose. The two most near and dear are the telephone network and computer time sharing. Other utilities—basic services provided by a regulated, quasi-monopoly—offer instructive models as well.

In the case of telephony we have a technology that is fundamentally *circuit-switched*—a fixed bandwidth and route is allocated to a given connection. Measurement of resource utilization is quite straightforward in the telephone network in contrast to an internet where packet switched technology does not guarantee bandwidth or route for a connection. Measurement or call accounting is so straightforward in the world of telephony that there is virtually no literature on the topic. The other suggestive features of the telephone network are the distinction between local and long-distance calling, and local service that is predominantly priced for unlimited local access. Last but not least, information services and assistance, formerly free, have become extra-cost items in the wake of deregulation.

The other familiar model is resource accounting from the large-scale computers. In another era batch jobs and interactive processes were all accounted for and individual users would be presented with monthly bills summarizing their use of a number of computational metrics: CPU time, disk storage, elapsed time, connect time and so forth. From this aggregation of metrics computer centers would devise schemes to recover operational costs. At my college the great pastime in the economics and physics departments was to structure batch jobs that minimized cost for a given task (typically a SAS run). What is instructive in this case is the notion that there are several metrics that are pertinent to charging for computation and there is a great deal of institutional variation in what metrics are used to calculate billable activity. The analogy breaks down in that metrics for batch computing are well defined whereas appropriate metrics for internet resource accounting are not.

Other utilities offer reasonable models too. Joules and Therms offer reasonable analogues to packets. Flat-rate basic cable service and extra cost premium channels map reasonably onto pricing for quality of service. In short, we all have a store of everyday economic experiences that offer a first iteration as to how one might approach the issue of performing internet accounting and chargeback. But again, packet-switched technology is without technical precedent in any of the models that we examine. So it is a special case in several respects.

Finite resources

One thing to bear in mind is that the current state of affairs, a networking infrastructure heavily subsidized by the federal government with essentially no economic discipline, is as real a policy as any that may be subsequently proposed. The essentially feature of the current system is that there are no incentives to use the finite resources efficiently. It's a common good like the air we breathe, and if I help myself to a bit extra what's it to you? It's a fact of life that mechanisms of one sort or another always emerge to ration a finite resource. There is the golden rule: "he with the gold rules." There is the "I was here first." There is the "it's the governments obligation to take care of me" approach. The current discipline is untenable for the future.

continued on next page

Internet Resource Accounting *(continued)*

The larger debate centers on what is the right thing to do to foster the free interchange of ideas over a media that has inherent costs associated with it? In the near term a research agenda should focus on developing reasonable metrics of network utilization with which to implement subsequent policy.

Players and interested parties

There are several formal organizations addressing Internet resource measurement and utilization (notice how I avoid saying accounting—it seems to be one of those symbolically loaded words that draws ire from the "Live Free or Die" Internetters). And of course there are several organizations that have this interesting theoretical problem in spades *now*, not in the near future.

First there is ISO where the *OSI Reference Model* incorporates a *Management Framework* that is sufficiently general to cover the case of internet resource utilization quite handily. The framework defines a generalized accounting management activity which includes a closed loop system for calculating and informing users and providers of the usage of particular managed objects including the ability to enforce limits one the use of the managed object.

The *Internet Research Task Force* (IRTF) commissioned the *Autonomous Network Research Group* to take the lead in policy and economic consideration of the Internet. The tangible artifact to date is a well thought out overview of Internet resource utilization and feedback [2].

The IETF convened the *Internet Accounting Working Group* in May of 1990. The group's charter is to develop an "Internet Accounting Architecture," with the intent of producing a set of standards that will be practical/usable for experimentation in the near term. The group follows the model laid out in the OSI Reference Model. A set of three RFCs should be published by year's end.

Service providers everywhere have a growing interest in network resource accounting. The Regional Bell Operating Companies (RBOCs) with *Switched Multimegabit Data Service* (SMDS) tariffs in the wind though interested in principle can be expected to create very conservative pricing structures. Though I am not in any way close to any tariff-creating activity at any of the RBOCs, only the most primitive pricing practices will be introduced with SMDS or ISDN data services, things that are only an incremental advance/modification to fixed cost bandwidth. Knowledgeable readers should feel free to contradict.

A few institutions have implemented systems to allow for zero-sum cost recovery of their network operations. Notable among these are Washington University and BBN—on behalf of the MILNET. As federal subsidy for the regional networks abates, look for an increasing interest in this topic from service providers ("How are we gonna stay in business?") and consumers ("How do I know that I'm paying my fair share?") alike.

What's really going on

From the community of users the interest in usage monitoring is quite straightforward: current monitoring technology (dedicated monitors and SNMP-based management systems) is built for diagnostic, interventionist behavior. It is designed to measure network illness, not wellness. Long-term usage monitoring by contrast, is optimized for providing user feedback, measuring compliance with network policies, and providing a mechanism for rational cost allocation/recovery. It is fundamentally for measuring things in a healthy network.

So what's really going on is this: the IRTF has articulated a set of concerns that need to be addressed as the Internet is transmuted into a national information utility. The IETF has appropriated ISO terminology and conception as the architecture for collecting information that will be relevant to long-term network monitoring. There are several working groups within the IETF that require almost identical SNMP-MIB extensions (Remote Monitoring, Network Accounting and others). At the engineering level there is considerable debate as to what are useful techniques for collection and storing potentially large amounts of data. In one corner stands the "collect everything because you don't know what's important and retrospective analysis will reveal it—after all CPUs, disks and network bandwidth are all cheap, right?" group. In the other corner is the "statistical sampling is all you'll ever need" crowd. SNMP seems destined to play a role in data collection for resource accounting, but much of the information that one would collect is unstructured and large. Bulk-data transfer protocols will likely be more appropriate.

Effective long-term monitoring requires substantial ingenuity. The most significant problem is not the time-constrained processing of protocol frames, but rather the conflicting requirements for maximizing information about network activity, while at the same time minimizing long-term storage. Current commercial monitoring technology forces a complete sacrifice of one or the other. So at the very least, in considering the problem of long-term network usage and feedback, we will develop techniques that refine the art of network monitoring in general.

But as for usage sensitive billing on the Internet...well I won't be holding my breath for that one!

Annotated Bibliography

[1] Braden, Robert T., "A Pseudo-Machine for Packet Monitoring and Statistics," Proc. SIGCOMM '88, August, 1988, ACM, pp 200–208. *If you are considering working in this area, this is an important article.*

[2] Estrin, Deborah & Zhang, Lixia, "Design Considerations for Usage Accounting and Feedback in Internetworks," unpublished manuscript, July, 1990. *This is a well written overview. I am not aware that it has become a draft RFC, so its distribution is still limited.*

[3] Mills, Cyndi, Hirsh, Don, & Ruth, Greg, "Internet Accounting Background and Status," Internet-Draft, April, 1991. *The first of three documents that the Internet Accounting Working Group will publish. Reasonable discussion of design goals.*

[4] ISO/IEC JTC1/SC21 N4958, "Information Technology—Open Systems Interconnection—System Management—Accounting Meter Function," Seoul, May 1990, *and*

[5] ISO/IEC JTC1/SC21 N4971, "Information Technology—Open Systems Interconnection—Accounting Management Working Document—4th Version," Seoul, May 1990. *Two documents that provide a succinct view of OSI accounting.*

DON HIRSH is Marketing Director, OEM Products at Meridian Technology Corporation, a communication software company based in St. Louis, Missouri. Previously he was Associate Director at Washington University's Office of the Network Coordinator, where he developed his interest in network resource accounting after having to justify his budget to numerous deans and department chairs. He can be reached as: `hirsh@MeridianTC.com`.

[Ed.: Don't miss the Internet Accounting BOF at INTEROP 91!]

Long, Fat Pipes

by Carl Malamud

[Ed.: The following article is adapted from the book *Stacks—Interoperability in Today's Computer Networks*. The book will be published at INTEROP 91 in October].

Introduction

"But why would I want a gigabit to my desktop?" I had been enthusiastically describing the Gigabit Testbed projects coordinated by the Corporation for National Research Initiatives to a prominent member of the industry. Gigabits of network bandwidth, petabytes of secondary storage, and teraflop computers had all been bandied about in a description of tomorrow's high-speed networks.

"Why would I want a gigabit?" is similar to a question that was common a few years ago, "Why would a PC ever need a full 640 kbytes of memory?" Needless to say, as soon as people discovered spreadsheets, 640 kbytes was not only reasonable, it became the minimal acceptable amount. When we learn to make do with what we have, we sometimes forget that the driving force is not our ability to make do with the existing technological base but the demand by users to get work done.

To see why we might want gigabit networks, let's start again with the lowly PC. If you want to do computer-generated real-time graphics, think of the VGA interface on the PC. A VGA screen has 640 x 480 bits with 256 simultaneous colors. To support 256 simultaneous colors, you need one byte per pixel. If you are operating at 30 screens per second, generally recognized as the minimum acceptable rate for real-time video, you are generating a data rate of 73.728 Mbps.

Now extend this analysis to workstations. Consider screens with 1000 x 1000 pixels and at least 2 bytes per pixel (yielding 64,000 simultaneous colors). All of a sudden we have increased our data rate from 73.728 Mbps to 480 Mbps.

Not every user needs 480 megabits per second on a transcontinental basis. Not every user is going to need all this bandwidth at all times. However, a few users will need this bandwidth at a time, and there are many, many users on real networks. We need both the capacity to deliver this bandwidth to the individual as well as the aggregate bandwidth to handle large numbers of users.

Another way to see the demand for high-speed networks is to examine how other portions of the computing environment are growing. A high-level (but not unusual) personal workstation has 100 Mbytes of storage, operates at 1–20 MIPS, and has 4–8 Mbytes of main memory. A rule that has held true for many years is that the shared computer of today ends up being the personal computer of tomorrow. A 1–20 MIPS machine with 8 Megabytes of Memory a few years ago would have been a large, shared VAX, but is now a personal workstation.

To see what the personal computer of tomorrow will be like, look at today's larger systems. It is not at all unusual to see 1 Gbyte of secondary storage, 20–50 MIPS systems, and 32–128 Mbytes of main memory. In fact, it is now possible (though fairly expensive) to buy personal workstations in this range. Over time, workstations will start to reach these levels for large numbers of people.

Let us look at even larger systems. Supercomputer centers and research laboratories are already working with terabytes of data. Groups like NASA are beginning to think in terms of petabytes (thousands of terabytes) of secondary and tertiary storage and some people are beginning to think in terms of exabytes (millions of terabytes).

Current large scale processors operate in the billion operations per second range. The *High Performance Computing and Communications* (HPCC) initiative will pour serious money into the development of a computer that will operate at a trillion operations per second.

Finally, look at main memory. Large supercomputer installations like the NASA-Ames Research Center have systems with main memories in the gigabyte range. In addition to a gigabyte of main memory, these systems often have another gigabyte or two allocated as a RAM disk. Systems with a terabyte or more of main memory are not that far off.

Balance is the key to any computer configuration. If the machines are faster and the disk drives larger, the networks also need to grow. If you need to load data into a terabyte of main memory, you are going to need more than an Ethernet.

The rationale for high-speed networks is particularly compelling if you realize that certain computer facilities will not be able to be duplicated. Computers are expensive, particularly supercomputers. In many cases, it won't make sense to buy one of each for every site. Instead, we need to put different computers in different locations.

Users will need to put these disparate computing locations together to form solutions to problems. Many efforts are now underway that examine how different computing environments can be joined together to solve specific problems. In this article, we will look at two of these efforts.

Both the networks examined in this article, *VISTAnet* and *CASA*, are part of the National Gigabit Testbed program, coordinated by CNRI, the Corporation for National Research Initiatives. There are three other testbeds as part of the project. Funding for this program of $15.8 million comes from DARPA and NSF, with another $100 million in facilities, equipment, and personnel thrown in by a list of industrial participants that includes almost every major computer and telecommunications company in the country.

VISTAnet

The VISTAnet project is coordinated by the MCNC, a non-profit corporation that runs the North Carolina Supercomputing Center. The group also runs another network called CONCERT which is a statewide private network built strictly on microwave towers. The network consists of three video networks that deliver NTSC signals to classrooms for classes and teleconferencing, plus a data network.

Computers

The VISTAnet project brings together three types of computers. First (of course) there is a large *Cray* computer. In this case the Cray computer is a CRAY Y-MP 8/432 with four processors, 64 megawords of main memory and another 128 megawords in a solid state disk. The machine has a peak performance of 1.2 Gflops. The Cray computer is located in a supercomputer center in the Research Triangle Park in North Carolina.

The second computer is a *Pixel-Planes 5,* an experimental machine developed at the Computer Science Department at the University of North Carolina. This machine does high-speed rendering of graphics. This autistic computer has the ability to do this one type of operation, and only this one type of operation, very fast. The Pixel-Planes is ideal for rendering polygonal images with lighting, shadows, and textures. The third machine is a *MasPar,* a commercial parallel processor, used for performing statistical manipulations.

continued on next page

Long, Fat Pipes *(continued)*

All three of these machines are combined to help feed a workstation used by the Department of Radiation Oncology at UNC. The machine uses a joy stick to allow the physician to control a 1280 x 512 pixel color display that shows a three-dimensional representation of radiation doses.

The VISTAnet project links machines in three locations: the North Carolina Supercomputing Center (NCSC) in Research Triangle Park, and two departments at the University of North Carolina, the Department of Radiation Oncology and the Computer Sciences Department.

Switch technology

The three customer locations are near two different central offices, each operated by a different telephone company. University of North Carolina links up to a Southern Bell office. Research Triangle Park is served by GTE.

The two central offices are linked together with a 2.4 Gbps OC-48 SONET line. A Fujitsu FETEX-150B-ISDN ATM switch is placed in the Southern Bell office and provides OC-12C links to the two UNC departments, each link operating at 622 Mbps.

The Fujitsu switch is the primary switch for the network. The link to the OC-48 line moves data over to the GTE office. There, a broadband circuit switch moves data on toward the Cray computer. The circuit switch (also known as a digital cross connect) can also be used for other applications such as video teleconferencing.

Because computers do not have a raw SONET link, another standard is used to move the data onto the computer systems. A HIPPI to ATM *Network Terminal Adaptor* (NTA) provides this function. The computers have a simple HIPPI interface to the NTA.

The role of the NTA is to take incoming ATM cells and present them to the HIPPI interface at 800 Mbps. The NTA has to perform rate adaption between the ATM rate of 622 Mbps and HIPPI's 800 Mbps. The NTA also provides the connection management function. Remember that the HIPPI interface allows communication with one device at a time, even though the ATM interface allows multiplexing of traffic. The NTA blocks calls until a virtual circuit is available to a remote HIPPI interface.

Why VISTAnet?

VISTAnet is in place mainly to do networking research. It sure helps, however, when you have a user. The user for VISTAnet is a fascinating experiment that applies high-speed networking to an area of medical practice known as *radiation oncology*.

When a person gets cancer, there are three ways to treat it. Surgery and chemotherapy are often used, but suffer from many drawbacks. A third approach is to use *radiation therapy*. Radiation therapy takes a cancer and kills it with a beam of radiation. The problem is that both normal and diseased cells get killed when exposed to radiation. Since the beam must pass through healthy tissue, a beam strong enough to kill the cancer will also kill healthy cells.

Luckily, the effect of radiation on cells is dependent on the dose. Small exposures do not hurt cells. If we split a radiation dose up into several beams that intersect at the diseased area, we can kill the diseased cells because they receive exposure to all beams. Healthy cells only receive exposure to a single beam and thus are able to survive.

Planning radiation treatment strategies starts with a CT scan of the diseased and surrounding areas. Because each cancer is unique—each has its own location, shape, and size—a doctor must develop an individual treatment plan for each situation that kills the diseased area without killing healthy cells.

Developing treatment plans operates in a very large parameter space with an infinite number of possible solutions. CT scans allow a two-dimensional view of the area. Analysis of a treatment plan shows the distribution of the dose over diseased and healthy areas within a single plane, but does not show the effect of the beams above and below the plane.

Problems

We thus have two problems. First, the number of possible solutions is very large. The number of beams, the locations of the beams, the position of the patient, and the type of shielding are just a few of the parameters. The complexity of the problem means only a few possible plans are tried.

The second problem is that the two-dimensional nature of the analysis means that only a portion of the effect can be seen. The result is that it is not uncommon for a treatment to get most of a diseased area, but not all, known as a *local failure.*

Estimates are that over 390,000 patients are treated per year with radiation therapy. Of those, 38,000 of the treatments are subject to local failures. While better diagnosis and treatment plans would not save all of those patients, it is likely that several thousand lives could be saved with better treatment.

Radiation oncology is thus a great application for high speed networking. When the VISTAnet project was trying to find an application for a gigabit testbed, they approached Dr. Julian Rosenman at the University of North Carolina.

Rosenman explained his problem in radiation oncology and his large computational requirements. To calculate a radiation dose distribution, a model consists of a system of 256^3 points. Each of these data points occupies 16 bits per point, resulting in a data rate of 256 megabits, clearly the province of a Cray computer.

Graphic representation

The information coming out at this rate is not very useful as raw data. It needs to be graphically represented to be useful to a physician looking at alternative treatment plans. The data stream goes into the Pixel-Planes 5 machine, 18 miles away from the Cray computer. This system is able to render the incoming data stream in near real time, at which point it needs to be displayed on a workstation, resulting in another very large data stream going over to the workstation.

VISTAnet is an example of how a single user can easily use a gigabit network. Visualization of radiation doses is an application that could not work in one site. It requires scarce facilities in multiple locations.

Note that in the medical field, it is not just the supercomputers that are scarce resources. Medical equipment is often very expensive and cannot be duplicated. Networks allow this scarce medical equipment to be used along with other scarce resources in other locations to form a more complete picture of diagnosis or treatment.

continued on next page

Long, Fat Pipes *(continued)*

CASA

A second gigabit testbed project is CASA. CASA involves four of the most highly developed computing centers in the country:

- San Diego Supercomputer Center (SDSC)
- California Institute of Technology (Caltech)
- Jet Propulsion Laboratory (JPL)
- Los Alamos National Laboratory (LANL)

All four of these sites are known for providing the latest in supercomputer facilities. Caltech and LANL are both leaders in applying parallel processors such as the Connection Machine to real-world problems.

With all these large facilities, however, there are a series of problems that can overwhelm any one of the computer centers. CASA will tie all four of the sites together with an 800 Mbps computer network, spanning up to 1300 kilometers.

CASA involves three applications, two of which are described in this article, that require very high-speed networks. Like the VISTAnet radiation oncology example, they are real-world problems that require solutions unavailable on any one large computer system, or even in any one large computer center.

Predicting the weather

Weather modeling is one of the applications that helps to spur larger and larger computers. Our weather system is so complex that most models concentrate on either the ocean or the atmosphere. Even dividing the problem into two leads to immense amounts of data.

Take atmospheric models, for example. If we take the world and divide it into "squares" of five degrees longitude by four degrees latitude, we have a fairly coarse grid of the world. If we model nine altitude layers, we have a grid of 72 x 44 x 9.

Even this coarse model of the atmosphere requires ten CPU seconds on a CRAY X-MP/48 to advance the model one hour. If we want to study a particular weather phenomenon, such as the Greenhouse Effect, it is not unusual to run a model through 50 years of space, requiring about 35 CPU days on the Cray computer.

Remember, this simplified model represents only one-half of the weather system, the atmosphere. The ocean model, at a coarse approximation, is a grid of 360 x 180, 27 levels deep. To advance this model a single hour, takes 20 CPU seconds on the CRAY X-MP/48, twice as long as the atmospheric model.

Although the atmosphere and the ocean have been separated, the weather should really be treated as a closed system. The output from the ocean model, especially sea surface temperature, is a key input to the atmospheric model. Outputs from the atmospheric model, such as winds and heat flux, are key inputs to the ocean model.

Under the direction of R. Mechoso of UCLA, CASA will combine two standard models to form a single closed system, the input from one model driving the other. Two machines will be used, one for each model.

The oceanic model will be put on a Connection Machines CM-2, which speeds the ocean model up by a factor of 50–100 times. The speedup is due to the superior programming model, for this particular application, of the massively parallel architecture.

The atmospheric model will be put on a CRAY Y-MP 8/864, located at the San Diego Supercomputer Center. This particular configuration of the CRAY Y-MP yields speedup of twelve times over the CRAY X-MP. Notice that the ocean model will be running fifty times faster, whereas the atmospheric model will run only a dozen times faster.

To handle the mismatch, the hydrodynamic part of the atmospheric model will be moved over to the CM-2, leaving the Cray computer with atmospheric problems like cumulus cloud convection and radiation calculations. Of course, a tremendous amount of data needs to move between the Cray computer and the CM-2, including data such as temperature and humidity for each grid for each cycle. Estimates are that approximately 750 Mbps per second of data will be transferred between the two machines.

Why bother with all this? A unified model is a way of tuning individual components so they reflect reality much more closely. If valid inputs yield valid outputs, we can start looking more accurately at questions like the Greenhouse Effect, forecasts of trade winds, and other global phenomena.

3-D Seismic profiling

Before we look at the CASA network itself, we examine one more CASA application, involving three-dimensional rendering of data from multiple earth-science data sets. This project is run at the Jet Propulsion Laboratory, but takes advantage of data and computers in different CASA locations.

Data in the earth sciences is increasing at fairly astonishing rates. There are a variety of different sources of this information: LANDSAT, topographic databases, and seismic databases, for example. One of the real challenging sources of information will be the space station, known as the NASA *Earth Orbiting System* (EOS). The EOS will be sending data down at the rate of 300 Mbps, equivalent to ten Gbytes every six minutes. And this is just one of many sources.

Combining information from different sources allows a variety of very important applications, including the modeling of earthquake faults, which allows prediction of an estimate of the order of magnitude of a coming earthquake (but not the exact time).

Earth sciences databases can be used for a variety of other tasks. Combined data sets have allowed researchers to discover that the Sahara desert was once a large river basin and even to find long-hidden roadways in Mongolia and Arabia, buried for several thousand years.

The point of the CASA application is to try and learn how to handle these very large datasets coming from different locations. For the JPL, this project is preparation for the flood of data expected from the space station. JPL is trying to learn how to handle data streams of three Gbytes/second and up which could require 90 Gigaflops or more to process.

Being able to handle data quickly is often crucial. An example is when the Voyager-2 was approaching Neptune. When the Voyager was 3–4 days out from the closest approach, an interesting feature was found on Neptune. Normally, it would take VAX systems weeks to analyze the data and provide positioning instructions for the on-board cameras. Instead, an eight-node Mark IIIfp was used to make the calculation quickly enough to send up repositioning instructions.

continued on next page

Long, Fat Pipes *(continued)*

The particular application chosen will merge data from three sources to provide 3-D cutaways of the earth's surface, allowing the identification of fault zones and major plate thrusts. Interactive 3-D graphics are essential for this application, because researchers cannot tell ahead of time the level of detail and particular view they need when examining specific places in the earth.

The three sources of information include the LANDSAT thematic mapper, CALCRUST seismic reflection data, and elevation data from the Space Shuttle's imaging radar. The amount of data involved for each image produced is fairly amazing.

Filtering

The LANDSAT thematic mapper, for example, involves a typical image of 90 x 90 kilometers. The image is broken up into 3000 by 3000 pixels with seven bands at ten bits, yielding 82 megabytes per data image. The shuttle elevation data form a 200Mbyte raw data set that needs to be filtered each time to yield a 6000 x 6000 point image. The seismic database is 1–2 gigabytes, taking tens of hours on a VAX to reduce to the pertinent information needed for a single image.

Once the three databases have been filtered, it takes yet more computer power to combine them to yield a rendered image. On a VAX, for example, it takes 14–17 minutes per frame for rendering. A minimal animation would be 1400 frames, requiring over 16 days of computing time.

The strategy to solve this problem is to break the problem down. Rendering of the data is performed on JPL's CRAY X-MP/18. The actual data filtering is done at Caltech, SDSC, and LANL. Figure 1 shows the extent of the data filtering. Even with the processing done at remote sites, there is still, if you have only one frame per minute, roughly 800 Mbps of data flow to the JPL.

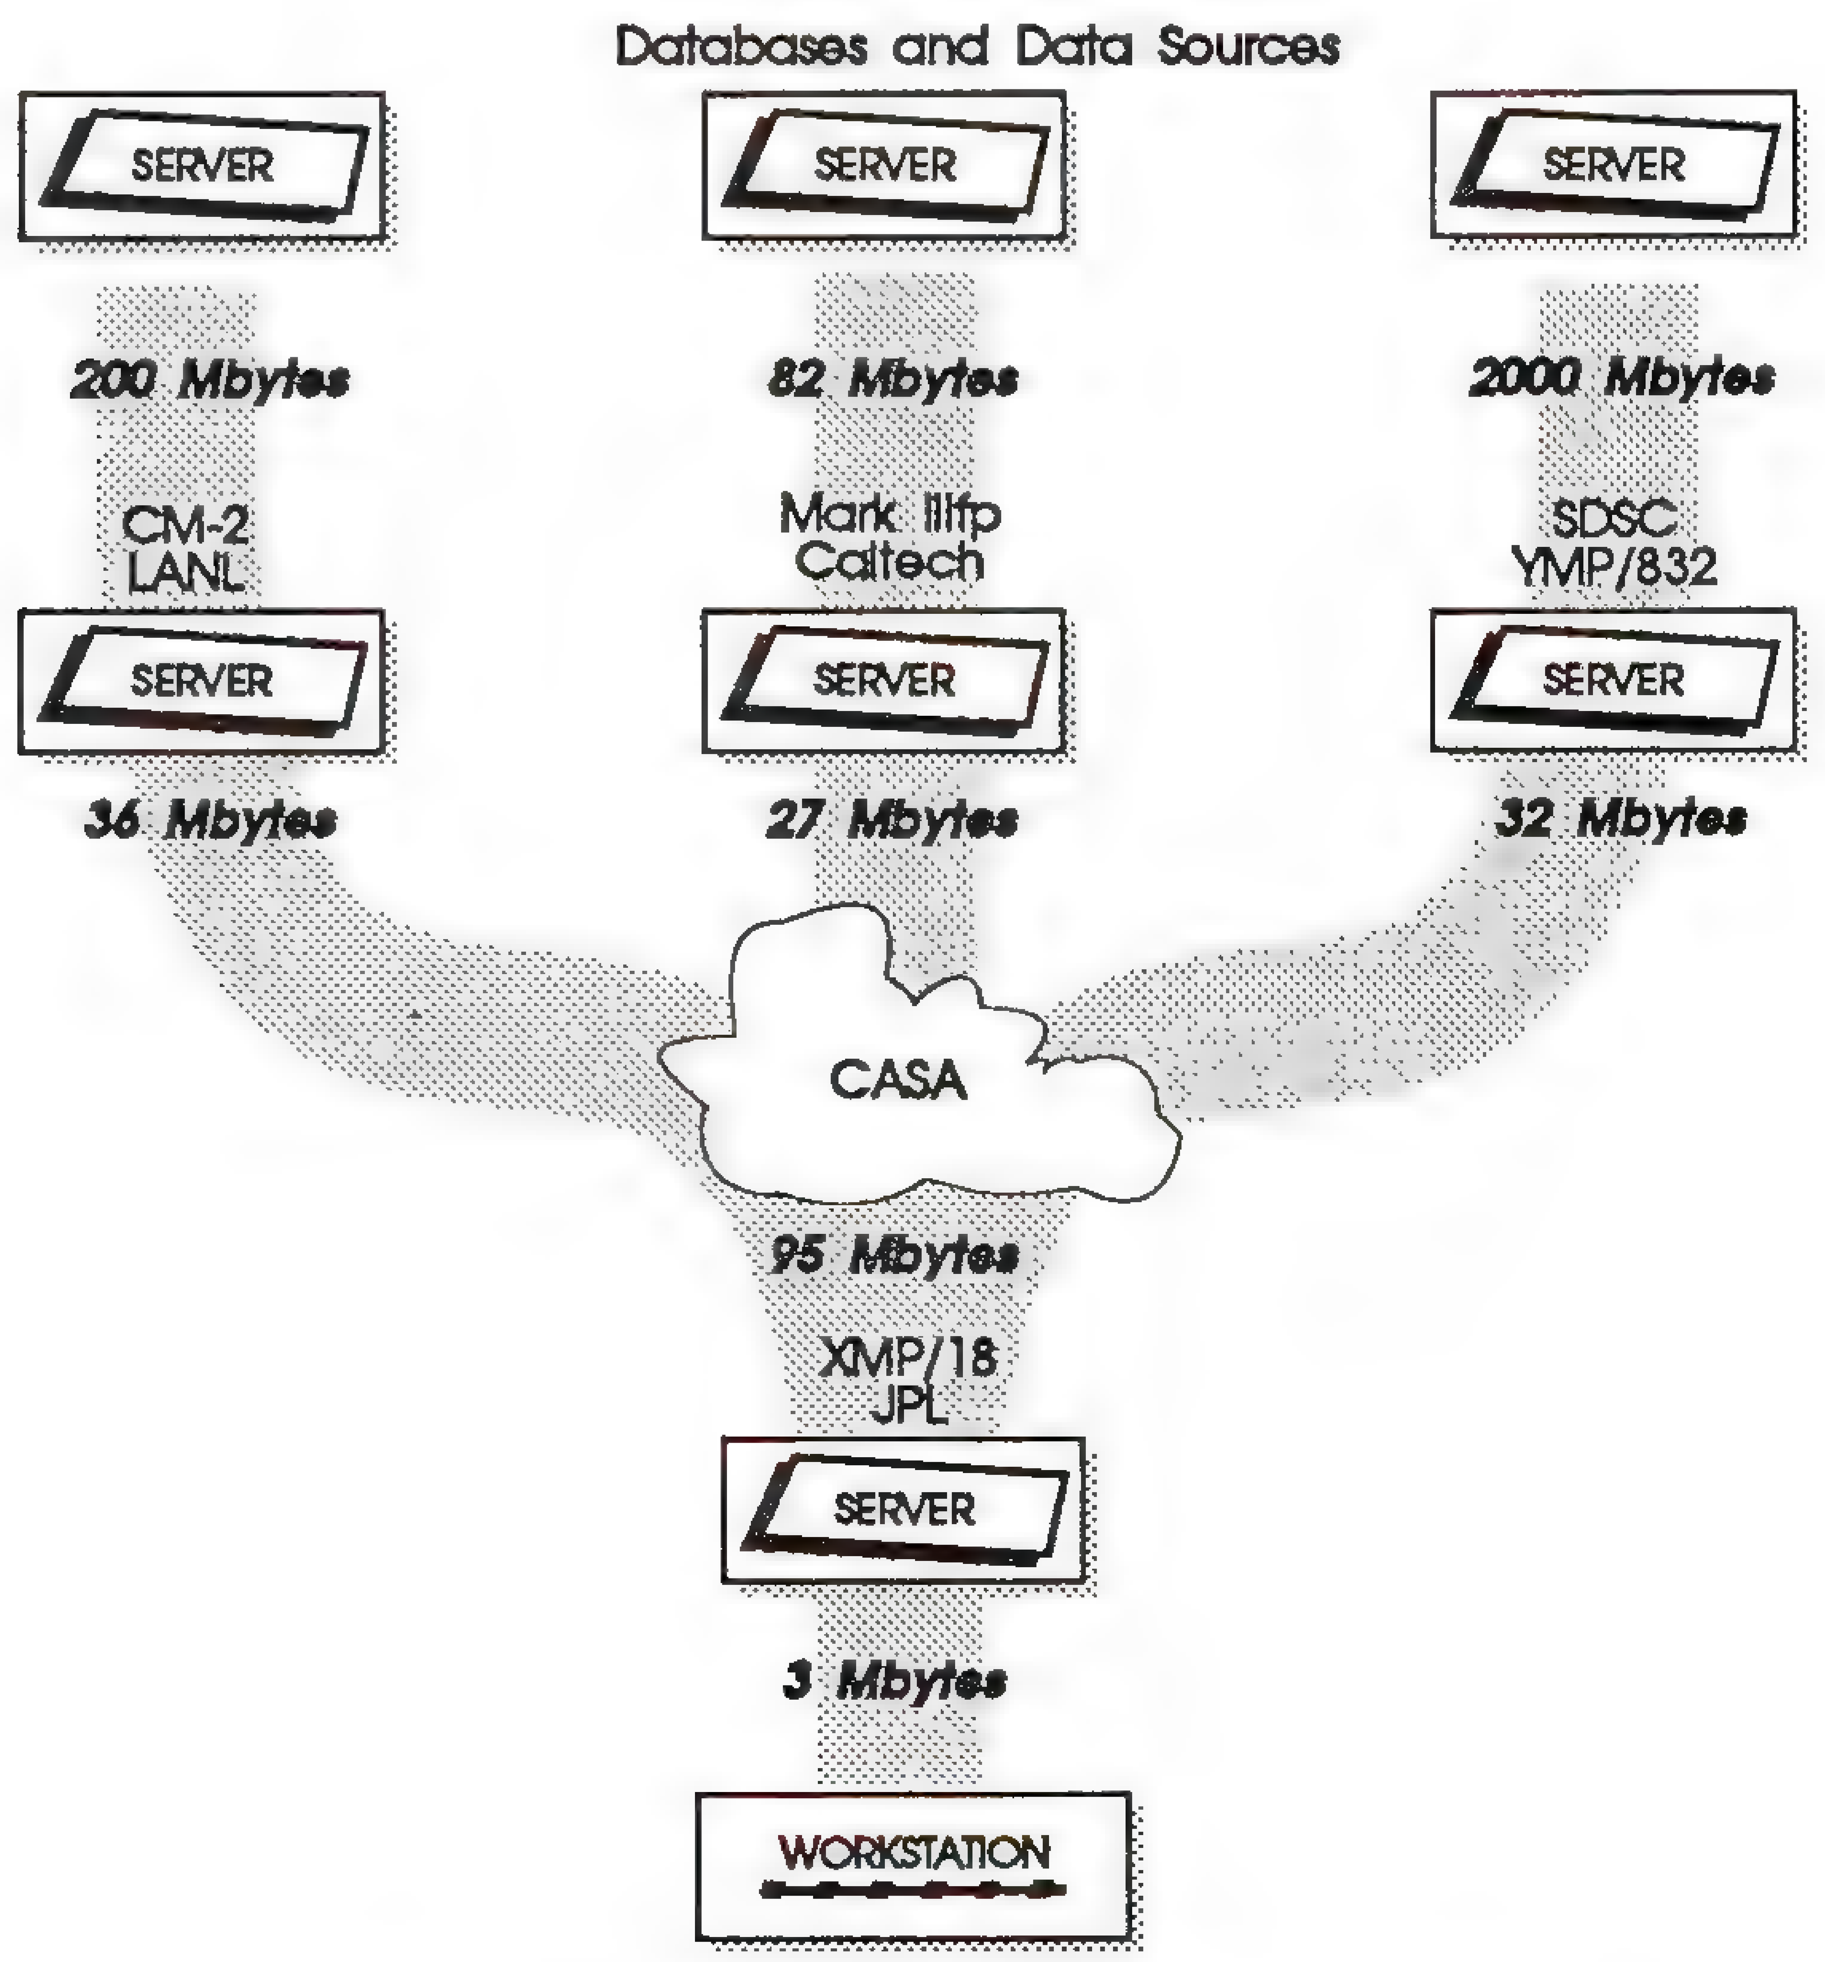

Figure 1: Dataflow in the CASA network

If you wanted to do real animation, it would take a minimum of 30 frames per second, yielding a data rate of 23.5 Gbps. The amount of processing power to do this animation would be 63 Gigaflops (the four machines involved in the application deliver around two Gigaflops of processing power).

The CASA network

The CASA network is a wide-area network. Each of the participating sites has a high-speed LAN, based on HIPPI. The host computers are all connected to the HIPPI switch. A HIPPI-SONET gateway is connected to the HIPPI switch.

The HIPPI-SONET gateway hooks up to long-haul optical fiber running the SONET protocols at STS-24 speeds (1.244 Gbps). Notice that SONET is being used directly instead of using an intervening ATM-based data link.

Linking HIPPI to SONET poses at least two problems. First, there is a difference in speed, with HIPPI running at 800 Mbps. Aside from rate adaptation, there is the more crucial problem of hiding latency. HIPPI won't let a source send data unless it has a ready signal. With a host required to store 64 ready HIPPI signals, and the propagation speed of HIPPI, we have a maximum HIPPI limit of 64 kilometers. CASA, however, needs 1500 to 2000 kilometers to function.

For HIPPI switches, CASA uses a switch developed at LANL in collaboration with DEC. The switch is a physical cross-bar switch; the switch actually moves to make the connection. This is a very fast physical switch, however, allowing a connection to be made in five microseconds if there is no contention.

The fiber for CASA is furnished by three telephone companies: MCI for the long-haul portion and the relevant Bell Operating Companies for the local loops. Built on top of this substrate is, to begin with, straight TCP/IP. If TCP proves inadequate as a transport layer, other candidates such as VMTP and NETBLT might be tried.

None of this, however, will be seen by the application programmer. The programmer would see, at the very lowest level, the UNIX sockets interface to TCP. Most programmers would work at even higher levels, using a collection of library routines such as *Express*. Express was designed for embedding information in C or FORTRAN programs that run on massively parallel processors. Express handles the questions of moving data and messages around and includes a symbolic parallel debugger and performance monitor for testing applications.

Serious?

Is the project serious? In addition to tremendous manpower, it is interesting to look at how much CPU time has been allocated on the big machines:

- SDSC has allocated 1900 CPU hours on the Y-MP 8/864.
- LANL has allocated 1100 hours on their Cray computer and 1100 hours on the CM-2.

The Cray computer CPU hours are supplemented by numerous other computers, not to mention a 1.2-Gbps, 1300-km fiber line.

CARL MALAMUD is the author of *DEC Networks and Architectures* (McGraw-Hill, 1989), *INGRES* (Van Nostrand Reinhold, 1989), *Analyzing Novell Networks* (Van Nostrand Reinhold, 1990), and *Analyzing DECnet/OSI Phase V* (Van Nostrand Reinhold, 1991). He has just finished writing *STACKS*, the official book of INTEROP 91, from which this article is adapted. The official caffeine-free diet soda of INTEROP 91 has not yet been chosen. Carl can be reached as carl@malamud.com.

Book Review

Practical UNIX Security, by Simson Garfinkel and Gene Spafford, O'Reilly & Associates, Inc., ISBN 0-937175-72-2, 1991.

Computers running variants of the UNIX operating system have achieved widespread use world-wide. However, given UNIX's modest upbringing in the academic and research communities, many of today's UNIX-based systems are vulnerable to numerous security threats. Everyone knows why it is important to safeguard the computers in one's care. What has been sadly lacking is a useful set of guidelines explaining what the threats are in a UNIX environment and how to prevent them or respond to them. As such, I am quite pleased by the publication of *Practical UNIX Security*. Indeed, the only thing that doesn't please me about this book is that in writing a review, there are so very few things that I can complain about!

Garfinkel and Spafford have collaborated to produce a book that every UNIX system administrator should read, and read again. In under 500 pages, the authors manage to concisely describe the spectrum of security concerns for non-military uses of UNIX. Rather than limiting themselves to just one variant of the UNIX operating system, their style of exposition is to describe the concepts in a UNIX-generic fashion and then delve into considerations specific for a given UNIX variant, such as BSD UNIX or System V UNIX.

Organization

The text is organized into five parts. The first two parts deal with the basics of the UNIX security model (users, groups, and the filesystem), and how to properly administer a system (e.g., restricting privilege, examining audit files and so on). These two parts provide solid information for the person managing a UNIX system with several users, but not connected to the outside world. This leads to the third part of the text which deals with securing a system with either dialup or network connections. The discussion is often not for the faint of heart as the number of security problems inherent in some network services is truly legion. After reading this part, I suspect many will vastly restrict their use of networked file systems. To be fair though, the coverage of network services is quite broad, and few, if any, emerged unscathed. (Indeed, it is the coverage of UNIX security and network services that motivated the review of this book in *ConneXions.*)

The fourth part of the book concerns itself with what happens, and what you should do, after a security problem occurs. It concludes with a rather sobering synopsis of the legal options available. The tone here is rather depressing given the rather limited positive results and the numerous (and often paradoxical) negative outcomes. For example, if a break-in occurs, and an investigation results, your computers might be seized as evidence—depriving you of further use of your resources. Whilst in hindsight this may seem common-sensical, the point is that the decision to engage in legal action for a remedy should be taken only after diligent evaluation of the possible outcomes.

The fifth part of the book is somewhat more erudite, dealing with the somewhat obscure but still important topics of *encryption* (primarily DES and RSA) and physical security (for hardware and data).

Checklist

The best feature of the book is the combination of explanation and checklist which is found in Parts 2 and 3: the authors begin by outlining the basic capabilities made available and then describe the numerous safeguards which must be employed in order to guard against attack. (The checklists from all the chapters are summarized in appendix for ease of reference.)

The authors scrupulously never give step-by-step information on how to exploit a particular weakness, but they do explain how to determine if the weakness is present in your system, and what steps to take accordingly. Sadly, many times the advice indicates that "you'll need source or must rely on your UNIX vendor." This illustrates how dependent our community has become on the "vendor element." However, the book provides an excellent checklist that users can present to their vendors. In some small way, this may help to raise the level of quality associated with some off-the-shelf UNIX systems.

Kerberos

The book's one fault is that in trying to be concise, the authors occasionally slip up and do not present enough background material to the reader to allow understanding of the security or environment issues involved. For example, when *Kerberos* is introduced in Part Three, the authors don't really explain what Kerberos is—just the environment that motivated it, and a very terse description of what it's like to run it. Then you get over a page of "What's Wrong with Kerberos?" The problem here is that the explanation of Kerberos is relegated to an appendix, so the casual reader won't really be able to appreciate if Kerberos's weaknesses are important in a particular environment. However, given the scope of the book and the excellent treatment of topics by the authors, this is a minor criticism at best.

Recommended

So, to sum up: if you administer a UNIX system, particular one with network or dialup connectivity, then you want this book. Read it once, and then refer to it often.

—*Marshall T. Rose*

Components of OSI

Have you been following our series of articles *Components of OSI?* Well, it's not over yet. We still plan to publish tutorials on ASN.1, ODA, The X.400 Message Store, A programmer's perspective on X.25 and more. Make sure you don't miss any of these in-depth articles on this important emerging technology. Back issues can be ordered by calling 1-800-INTEROP or 415-941-3399. Ask for the *ConneXions* subscription department. Since the *Components of OSI* series started we have covered the following topics:

ISDN	April	1989
X.400 Message Handling System	May	1989
X.500 Directory Services	June	1989
The Transport Layer	July	1989
Routing overview	August	1989
IS-IS Intra-Domain Routing	August	1989
ES-IS Routing	August	1989
The Session Service	September	1989
CLNP	October	1989
The Presentation Layer	November	1989
A taxonomy of the players	December	1989
The Application Layer Structure	January	1990
FTAM	April	1990
The Security Architecture	August	1990
Group Communication	September	1990
X.25	December	1990
The Virtual Terminal ASE	January	1991
Systems Management	April	1991
CO/CL Interworking	May	1991

Letters to the Editor

Ole,

The excellent review in the May issue of *ConneXions* of "Computers at Risk" listed only the 800 number for telephone orders. This number is unreachable from the vicinity of National Academy Press in Washington, DC. From within the 202 area code, the same destination as 800-624-6242 is reached by 334-3313.

John Schnizlein
Senior Communication Specialist
House Information Systems
U. S. House of Representatives

Dear Editor,

The May 1991, *ConneXions* with Mike Schwartz's article on "Resource Discovery..." came just as I had responded to a user's request to help locate people and hosts at an international site to be visited for several months as part of a research project. I had been "discovered" by word-of-mouth and I, in turn, passed the request on to people who were closer to the specific information sought: a people-based discovery model. Because you were thoughtful enough to have included Mike's Internet address with the article, 'twas but a twinkling of time 'til my query to Mike about the usefulness of *netfind* in these instances brought the following reply:

> "*Netfind* itself wouldn't help much for this particular search, because *netfind* focuses on finding individual users, rather than network connections. The techniques behind it could be used to solve your problem, though.
>
> It seems there are two issues here: discovering the resources, and getting permission to use them. I know of several ways to discover networks and machines in various parts of the world (including domain zone transfers and the *netfind* seed database, for example), but that still leaves the permission problem. I suppose one could send a message to `postmaster@discovered_machine`to get permission, but that seems like an unlikely avenue for success. Personal networks are probably the best for getting permission.
>
> On the other hand, the situation would be much different if there were a global distributed service that registered such information. I am currently drafting the plans for an Internet Research Task Force-chartered research group on resource discovery. Internet connections for people on the road might be an interesting application to consider.
>
> I think the whole game will become more complicated and interesting as the Internet commercializes, because then there might be competitive service providers."

It is encouraging to see this line of investigation and development taking place. Putting end-users within convenient reach (including the discovery phase) of the resources on the net is vital to achieving the promise of computer networking for all communities served. And, putting the networking community in touch with timely developments is a much appreciated service of *ConneXions*. Hats off to you, Ole.

Steve Goldstein
Program Director
National Science Foundation
Division of Networking and Communications
Research & Infrastructure

Upcoming Events

Themes

The *National Conference on Computing and Values* (NCCV) will convene August 12–16, 1991, in New Haven, Connecticut. NCCV 91 is a project of the National Science Foundation and the Research Center on Computing and Society. Specific themes (tracks) include:

- Computer Privacy & Confidentiality
- Computer Security & Crime
- Ownership of Software & Intellectual Property
- Equity & Access to Computing Resources
- Teaching Computing & Values
- Policy Issues in the Campus Computing Environment

The workshop structure of the conference limits participation to approximately 400 registrants.

Speakers

Confirmed speakers include: Ronald E. Anderson, Daniel Appleman, John Perry Barlow, Tora Bikson, Della Bonnette, Leslie Burkholder, Terrell Ward Bynum, David Carey, Jacques N. Catudal, Gary Chapman, Marvin Croy, Charles E. M. Dunlop, Batya Friedman, Donald Gotterbarn, Barbara Heinisch, Deborah Johnson, Mitch Kapor, John Ladd, Marianne LaFrance, Ann-Marie Lancaster, Doris Lidtke, Walter Maner, Diane Martin, Keith Miller, James H. Moor, William Hugh Murray, Peter Neumann, George Nicholson, Helen Nissenbaum, Judith Perolle, Amy Rubin, Sanford Sherizen, John Snapper, Richard Stallman, T. C. Ting, Willis Ware, Terry Winograd, and Richard A. Wright.

More information

The registration fee is low and deeply discounted air fares are available into New Haven. To request a registration packet, please send your name, your e-mail *and* paper mail addresses to:

Professor Walter Maner
Department of Computer Science
Bowling Green State University
Bowling Green, OH 43403
USA
BITNET: `MANER@BGSUOPIE.BITNET`
Internet: `maner@andy.bgsu.edu`
Fax: +1 419-372-8061
Phone: +1 419-372-8719 (answering machine)
+1 419-372-2337 (secretary)

Write to *ConneXions*!

Have a question about your subscription? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Want to enquire about back issues? (there are now more than fifty to choose from; ask for our free 1987–1991 index sheets). We want to hear from you. Simply write, call, fax, or e-mail to:

ConneXions—The Interoperability Report
480 San Antonio Road, Suite 100
Mountain View, CA 94040-1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: `connexions@interop.com`

I look forward to hearing from you!

Ole J. Jacobsen

Ole J. Jacobsen
Editor & Publisher

Call for Papers

The *Eleventh Annual IEEE Phoenix Conference on Computers and Communications* (IPCCC-92) will be held April 1–3, 1992 in Scottsdale, Arizona. This international conference provides a forum for the presentation and exchange of current work in the field of synergism of computers and communications, and their applications areas. We are particularly soliciting industrial, business, and government participation as well as the active involvement of the academic community. We know it is vital that there be a dialogue between practitioners and researchers. Thus, in addition to research contributions, we look forward to reports detailing experiments, evaluations, problems, and opportunities associated with design, implementation, and operation. Such reports will be given special consideration.

Papers

Submitted manuscript must be no longer than 5,000 words, be typed double-spaced, and include an abstract of approximately 300 words. Long papers and reports will not be considered. Authors should obtain company and government clearances prior to submission of the papers. Please submit five copies of complete paper and abstract by 15 July 1991 to:

Dr. Ming T. (Mike) Liu, Program Chair
Ohio State University
Department of Computer & Information Science
2036 Neil Avenue
Columbus OH 43210-1277
Voice: 614-292-6552
Fax: 614-292-9021
E-mail: `mike.liu@osu.edu`

All papers submitted will be refereed by the Program Committee. They will be judged with respect to their quality originality and relevance. Authors will be notified of acceptance/rejection shortly after 20 September 1991. Accepted papers will be published in the IPCCC-92 Proceedings.

Sessions and Tutorials

We solicit proposals for special topics and panel sessions. Each proposal should include subject, justification, and names of possible participants. Proposals should be sent to the Program Chair by 15 July 1991. Proposers will be notified of provisional acceptance shortly after 15 August 1991. Proposals for one-day tutorials related to the suggested topics are also desired. Please contact the Tutorials Chair for tutorial proposal submission guidelines. Proposals should be sent by 15 July 1991 to:

Dr. Ann Miller, Tutorials Chair
Motorola Satellite Communications
2501 South Price Road, Mail Drop G1140
Chandler, AZ 85248-2899
Voice: 602-732-3874
Fax: 602-732-3046
E-mail: `am-ipccc@asuvax.eas.asu.edu`

Exhibits

Exhibits of commercial products and demonstrable prototypes related to the suggested topics are solicited. Please contact the Exhibits Chair by 3 September 1991:

Dr. Frank Caliss, Exhibits Chair
Arizona State University
Department of Computer Science & Engineering
Tempe, AZ 85287-5406
Voice: 602-965-2804
E-mail: `caliss@asuvax:eas.asu.edu`

Suggested topics

Topics within the scope of the conference include, but are not limited to:

- Parallel and Distributed Computing
- Fault Tolerance and Reliability
- Neural Network Computing
- Distributed Database Systems
- Optical Disk Storage
- VLSI/VHSIC Developments
- Advanced Architectures
- Fiber Optics
- Satellite/Terrestrial Systems
- Communications Theory
- Spread Spectrum
- Specification Methodologies
- Development Environments
- Object-Oriented Systems
- Real-Time Systems
- Performance Measurement and Evaluation
- Graphics and Scientific Visualization
- Distributed Operating Systems
- Project Management
- OSI Networks and Interoperability
- Fault Tolerance Networks
- Local and Wide Area Networks
- Network Management, Control, and Security
- ISDN Systems
- Value Added Networking
- Expert System Design and Applications
- Non-traditional Languages
- Distributed AI Systems
- Intelligent Databases
- Medical Information Systems
- Process Control
- CAD/CAE/CAM
- Robotics and Computer Vision
- Multi-Media Databases

Preliminary Announcement and Call for Papers

Background

The USENIX *Mach Symposium* will be held at the Doubletree Hotel in Monterey, CA, November 20–22, 1991. Mach has become a dynamic addition to the operating systems marketplace. DARPA originally sponsored Mach development, and continues to emphasize the use and growth of Mach. In the larger research community, Mach is ever more widely used at many university sites and industrial research labs. Versions of Mach have been released commercially by Encore, NeXT, BBN and Mt. Xinu. The *Open Software Foundation* (OSF) chose Mach as the basis for its operating system offering, OSF/1. Mach is finding increasing acceptance as computer vendors ready products derived from OSF/1.

Recent developments have demonstrated the feasibility of Mach 3.0, the combination of a pure Mach kernel with single or multiple servers emulating the features of traditional operating systems. Performance of Mach 3.0 has begun to approach or exceed that of Mach 2.5. Workers outside of the CMU community have begun to use Mach 3.0 as the basis for their projects. In short, acceptance of Mach has come about in an astonishingly brief time.

Activity in this field has been sufficiently wide-spread that, little more than a year after the first USENIX Mach workshop, the USENIX Association is pleased to sponsor an expanded Mach symposium to bring together researchers, engineers, vendors and users of Mach systems. We will encourage discussion of all past and present Mach-related research, development, production and applications activities.

Symposium overview

The symposium will be spread over three days. The first day will be devoted to two half-day tutorials on advanced programming for Mach 3.0. The following two days will concentrate on presentation of refereed papers on current and historical Mach-related work. Long breaks between presentations provide ample opportunity for informal discussion. Some time will be available for descriptions of work in progress.

Tutorials

Richard Draves will lead a tutorial entitled *Writing a Multi-Threaded Mach 3.0 Server,* with particular attention paid to the complexities of using Mach IPC. During the course of his doctoral studies at Carnegie-Mellon University, Rich rewrote Mach 3.0 IPC to solve problems that became apparent with Mach 2.5 servers.

David Black's tutorial is called *Writing an External Memory Manager.* Discussion will center on the intricacies of developing an efficient (and well-behaved!) external manager. David, currently of the Open Software Foundation, received his doctorate from Carnegie-Mellon for his contributions to Mach.

These tutorials are being developed precisely for this USENIX Mach symposium. They will explore concepts and rationale as well as real examples. They are oriented towards programmers who already have some familiarity with using Mach IPC and VM. Each tutorial is a half-day, so conference attendees may take part in both. The tutorials will be priced separately from the conference registration fee.

Submissions

Extended abstracts of 1500–2500 words (9000–15000 bytes or 3–5 pages) should be sent to Alan Langerman at the address below. Shorter abstracts run a significant risk of rejection as there will be little on which we can base an opinion. Preference will be given to those submissions that include an outline of the entire paper in addition to the extended abstract. A good extended abstract will contain the following information in one form or another:

- Abstract: 100–300 words included verbatim in the final paper
- Introduction: The problem; its importance; previous work
- Solution: Rationale, decisions, tradeoffs. Implementation details.
- Evaluation: Performance results; effort required; lessons learned.
- Conclusion: Summary of the paper, explains its importance.

The extended abstract will allow us to analyze the content of your proposed paper. An outline lists the headings, major points and many minor points for each section of the actual paper. The submission package should include the paper title as well as:

- Your extended abstract
- Outline of rest of paper, if at all possible
- Estimate of paper length
- Contact author (liaison to program committee)
- E-mail address and daytime phone number for contact author
- Optional home phone number and optional FAX number
- Surface mail address (required)
- If you are submitting hardcopy, six copies of the submission

Topics

Areas of interest include, but certainly are not limited to:

- Applications
- Mach 2.5 and earlier development
- Mach 3.0 monolithic server
- Mach 3.0 multi-server
- Problems with Mach 2.5 / Mach 3.0 features
- Multiprocessor or parallelization experiences
- Security
- Performance
- Productization
- Experiences with OSF/1
- Use of Mach subsystems in other operating system kernels
- Comparisons of Mach with other operating systems; e.g., Chorus, Sprite, Amoeba, V, and of course UNIX
- Porting Mach to off-beat architectures
- Future work

Important dates

Extended abstracts:	July 19, 1991
Notification:	August 23, 1991
Camera-ready, full papers:	October 4, 1991

To submit a paper or for more information, contact the program chair:

Alan Langerman
Encore Computer Corporation
257 Cedar Hill Street
Marlborough, MA 01752
Voice: +1 508-460-0500
Fax: +1 508-485-0709
E-Mail: alan@encore.com

Call for Papers

The *USENIX Winter 1992 Technical Conference* will be held January 20–24, 1992 at the San Francisco Hilton in San Francisco, California.

Innovation

Some believe that UNIX standardization efforts have killed innovation. And yet, we need innovation. Large write-once disks make the current filesystem untenable. Even the 2 gigabyte file limit built in all through the system breaks. Gigabit networking clogs an I/O model designed to push hundreds of kilobytes per second, not hundreds of megabytes. System administration for thousands of machines? Programming tools for distributed workgroups? Object-oriented and visual programming? Microkernels with client/server architectures? RAID disk arrays? Transcontinental file servers? What's a programmer to do?

The USENIX Winter 1992 Conference solicits new work on all topics related to UNIX or UNIX-inspired systems programming and technology. But as always, we care most about innovation and how it coexists with (and sometimes thrives on) stasis. Please target a sophisticated technical audience particularly knowledgeable of operating system issues yet keenly interested in new and exciting projects in many areas.

Vendors are encouraged to submit technical presentations on products. However, we will reject obvious product announcements. Previously published papers will also be rejected, although "retrospective" papers may describe work done years ago.

Submissions

Submissions must be in the form of extended abstracts, 1500–2500 words in length (9000–15000 bytes or 3–5 pages). Shorter abstracts will not give the program committee enough information to judge your work fairly and, in most cases, this means your paper will be rejected. Longer abstracts and full papers simply cannot be read by the committee in the time available. However, you may append a full paper to an extended abstract; this is sometimes useful during evaluation.

The extended abstract should represent your paper in "short form." The committee will want to see that you have a real project, that you are familiar with other work in your area (i.e., include references), and that you can clearly explain yourself. You should have results and they should be summarized in your abstract. A good submission will contain:

- Abstract: This should be included verbatim in the final paper.
- Introduction:
  - Introduce the problem: why is it important?
  - Reference previous work.
- How We Solved the Problem:
  - More details on the problem and its issues.
  - Design decisions and tradeoffs, and why they were made.
  - Implementation details.
- Evaluation:
  - Data on performance and effort required.
  - How well does it work?
  - What would you do differently?
  - If it failed, why?
  - What did you learn from it?
- Conclusion:
  - Summarize the paper, emphasizing why it is important and what was learned.

In addition to the extended abstract, every submission should include:

- A clearly designated contact author who will be your link to the program committee.
- A daytime phone number (essential!).
- A surface mail address (required).
- An e-mail address, if available; e-mail is by far our best path of communication.
- A home phone number (optional, although questions often arise on evenings and weekends and it will avoid delays).
- A Fax number (optional).
- Any special audio/visual equipment you may require. A microphone, overhead projector, and 35mm projector will be provided as standard equipment. We are happy to provide additional assistance and equipment to make your presentation as audio and visually appealing as possible.
- Indication of student status.

Presentations are usually scheduled for 25 minutes.

Dates

The final date for submissions is August 19. Authors of accepted submissions will be notified by October 1. They will immediately receive instructions for the preparation of camera ready final papers to be published in the conference proceedings. Camera-ready papers of 8–12 typeset pages will be due by November 22.

Sending your paper

Submissions can be sent (in order of committee preference):

E-mail: `SFusenix@Usenix.ORG` or `uunet!usenix!SFusenix`

Paper mail: Eric Allman
Computer Science Division, EECS
University of California
Berkeley, CA 94720

Fax: +1 415-843-9461

Awards

A cash prize for the best paper by a full-time student will be awarded by the conference program committee. With your submission, please indicate if you are a full-time student. An Award for Best Paper at the conference is also made by the committee.

More information

Materials containing all details of the technical and tutorial program, conference registration, hotel and airline reservation information will be mailed in October 1991. If you did not receive a printed copy of this announcement directly and wish to receive the pre-registration materials, please contact:

USENIX Conference Office
22672 Lambert Street, Suite 613
El Toro, CA 92630
Phone: +1 714-588-8649
Fax: +1 714-588-9706

Announcing The Internet Society

by Vint Cerf, Corporation for National Research Initiatives

Background

The purpose of this article is to provide a brief description of the *Internet Society* and its goals and objectives. It will function as a professional society to facilitate, support and promote the evolution and growth of the Internet as a global research communications infrastructure. The suggestions and recommendations of all parties interested in the Internet are solicited to assist in making the Internet Society robust, productive and structured to meet the needs of its members.

The Internet Society

The Internet is a collection of cooperating, interconnected, multiprotocol networks which supports international collaboration among thousands of organizations. Because of its current scope and rapid rate of growth, the Internet will benefit from a more organized framework to support its objectives. To this end, an Internet Society is being formed to foster the voluntary interconnection of computer networks into a global research and development communications and information infrastructure. The Internet Society will *not* operate the Internet. Internet operation will continue to be a collaborative activity which the Society will seek to facilitate. The Society will provide assistance and support to groups and organizations involved in the use, operation and evolution of the Internet. It will provide support for forums in which technical and operational questions can be discussed and provide mechanisms through which interested parties can be informed and educated about the Internet, its function, use, operation and the interests of its constituents.

Membership

The Internet Society will be a membership organization with voting individual members and non-voting institutional members. There will be several classes of institutional members. The society will produce a newsletter on a regular basis and hold an annual meeting to which all members and other interested parties will be invited. The topics of the annual meeting will vary, but are expected to focus on current research in networking, Internet functionality and growth, and other interests of the Society constituency. All members will receive the newsletter and an invitation to attend the annual meeting of the Internet Society.

Membership dues will vary according to class of membership. The amounts of these dues and the basis on which they are set will be determined by the Board of Trustees of the Society and may be revised from time to time as set forth in the By-Laws.

Charter

The Society will be a non-profit organization and will be operated for international educational, charitable and scientific purposes among which are:

- To facilitate and support the technical evolution of the Internet as a research and education infrastructure and to stimulate involvement of the academic, scientific and engineering communities among others in the evolution of the Internet.
- To educate the academic and scientific communities and the public concerning the technology, use and application of the Internet.
- To promote scientific and educational applications of Internet technology for the benefit of educational institutions at all grade levels, industry and the public at large.
- To provide a forum for exploration of new Internet applications and to foster collaboration among organizations in their operation and use of the Internet.

Support for Internet Technical Evolution

The *Internet Activities Board* (IAB) has been concerned with the development and evolution of architectures supporting the use of multiple protocols in a networked environment. The Internet Society will incorporate the IAB and its functions into the operation of the Internet Society. The Internet Society will work with other interested organizations to support and assist efforts to evolve the multiprotocol Internet. The Society will use the Internet Engineering and Research Task Forces to stimulate networking research and facilitate the evolution of the TCP/IP protocol suite and the integration of new protocol suites (e.g., OSI) into the Internet architecture. The Internet Society will work with parties and organizations interested in fostering improvement in the utility of the Internet for its constituent users.

Meetings and Conferences

The Internet Society will convene an annual meeting and will organize and facilitate workshops and symposia, jointly with other organizations where appropriate, on specific topics of interest to the Society membership. The annual meeting will address issues of global and regional importance to the evolution and growth of the Internet. In particular, future INET conferences will be incorporated into the Society's annual meetings.

Information and Infrastructure Services

The Internet Society will publish an Internet Newsletter providing members with information about the international activities of Internet constituents. In addition, the Society will also provide assistance to and support for organizations responsible for maintaining the databases crucial to Internet function (e.g., the Domain Name System, X.500 Directory Services, etc.) and organizations concerned with the security of the Internet (e.g., the Software Engineering Institute Computer Emergency Response Team (CERT) and its CERT-System). The Society will assist in the development of educational, advisory and informative materials of use to Society members. Where appropriate, the Society will organize or support activities which aid in the coordination among the organizations operating components of the Internet.

The Society will refer members to appropriate parties involved in operating the various parts of the Internet where they may be helpful with specific questions. Where possible, the Society would seek to provide access to its information on-line, but would also offer hard copy and, perhaps eventually, CD-ROM-based information resources.

Plans

The initial organizers of the Internet Society include the Corporation for National Research Initiatives (CNRI), EDUCOM and the Internet Activities Board. During the six month period from June to December 1991, the initial organizers will work with interested parties to prepare for beginning operation of the Society by the end of 1991. Computer networking has become a critical infrastructure for the research and development community and has the potential to become the basis for world-wide collaboration and cooperation in every field of human endeavor. The Internet Society will seek to solidify, enhance and encourage further international collaborative networking. Visionary individuals joining the Society during its formation will receive special recognition as Society pioneers and will have the opportunity to shape the early agenda of Society activities. Opportunities for organizational and institutional participation are also available.

It is time. The technology is available. A global renaissance of scientific and technical cooperation is at hand. You are cordially invited to take part in an enterprise without precedent and an adventure without boundary. The Internet Society sets sail in January of 1992 on a voyage of internetwork discovery. Will you be aboard?

For more information, call CNRI at 703-620-8990.

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED